



US009246878B2

(12) **United States Patent**
Kumar et al.

(10) **Patent No.:** **US 9,246,878 B2**
(45) **Date of Patent:** ***Jan. 26, 2016**

(54) **METHODS AND SYSTEMS FOR ROUTING
PACKETS IN A
VPN-CLIENT-TO-VPN-CLIENT
CONNECTION VIA AN SSL/VPN NETWORK
APPLIANCE**

USPC 726/15
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

(71) Applicant: **Citrix Systems, Inc.**, Fort Lauderdale,
FL (US)

5,835,725 A 11/1998 Chiang et al.
6,049,826 A 4/2000 Beser

(Continued)

(72) Inventors: **Arkesh Kumar**, Santa Clara, CA (US);
James Harris, Santa Clara, CA (US);
Ajay Soni, Santa Clara, CA (US)

FOREIGN PATENT DOCUMENTS

(73) Assignee: **CITRIX SYSTEMS, INC.**, Fort
Lauderdale, FL (US)

WO WO-01/47185 A2 6/2001
WO WO-2006/012612 2/2006
WO WO-2006/074072 A3 7/2006

OTHER PUBLICATIONS

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 0 days.

Advisory Action for U.S. Appl. No. 11/462,174 dated May 4, 2010.

(Continued)

This patent is subject to a terminal dis-
claimer.

(21) Appl. No.: **14/045,922**

Primary Examiner — Ali Abyaneh

Assistant Examiner — Narcisco Victoria

(22) Filed: **Oct. 4, 2013**

(74) *Attorney, Agent, or Firm* — Foley & Lardner LLP;
Christopher J. McKenna; Paul M. H. Pua

(65) **Prior Publication Data**

US 2014/0041014 A1 Feb. 6, 2014

Related U.S. Application Data

(63) Continuation of application No. 11/462,174, filed on
Aug. 3, 2006, now Pat. No. 8,572,721.

(51) **Int. Cl.**

H04L 29/06 (2006.01)

H04L 12/46 (2006.01)

(52) **U.S. Cl.**

CPC **H04L 63/0272** (2013.01); **H04L 12/4641**
(2013.01); **H04L 63/166** (2013.01)


(58) **Field of Classification Search**

CPC H04L 63/0272; H04L 63/166; H04L
12/4641; H04L 63/08; H04L 63/20; H04L
63/164

(57) **ABSTRACT**

In a method and system for routing packets between clients, a packet is received from a first client connected to a secure sockets layer virtual private network (an SSL/VPN) network appliance. An identification is made, responsive to an inspection of the received packet, of i) a type of connection required for transmission of the received packet to a destination address identified by the received packet and ii) a second client connected via an SSL/VPN connection to the SSL/VPN network appliance and associated with the identified destination address. A request is made for establishment by the second client of a connection of the identified type within the SSL/VPN connection. The received packet is transmitted to the second client via the established connection of the identified type.

20 Claims, 10 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

6,151,631 A 11/2000 Ansell et al.
 6,189,102 B1 2/2001 Beser
 6,272,148 B1 8/2001 Takagi et al.
 6,324,582 B1 11/2001 Sridhar et al.
 6,353,614 B1 3/2002 Borella et al.
 6,381,646 B2 4/2002 Zhang et al.
 6,480,508 B1 11/2002 Mwikalo et al.
 6,490,289 B1 12/2002 Zhang et al.
 6,567,405 B1 5/2003 Borella et al.
 6,598,081 B1 7/2003 Coile et al.
 6,611,875 B1 8/2003 Chopra et al.
 6,614,809 B1 9/2003 Verma et al.
 6,631,416 B2 * 10/2003 Bendinelli et al. 709/227
 6,920,502 B2 7/2005 Araujo et al.
 6,983,319 B1 1/2006 Lu et al.
 7,072,933 B1 7/2006 Lamb et al.
 7,216,173 B2 5/2007 Clayton et al.
 7,240,150 B1 7/2007 Todd et al.
 7,246,373 B1 7/2007 Leung et al.
 7,386,881 B2 6/2008 Swander et al.
 7,403,535 B2 7/2008 Modi et al.
 7,450,560 B1 11/2008 Grabelsky et al.
 7,451,193 B1 11/2008 Coile et al.
 7,461,150 B1 12/2008 Boulia
 7,631,182 B1 * 12/2009 Droux et al. 713/151
 7,685,255 B2 3/2010 Amlekar et al.
 7,761,597 B2 7/2010 Takeda et al.
 7,769,869 B2 8/2010 Venkatraman et al.
 7,769,871 B2 8/2010 Boulia
 7,827,237 B2 11/2010 Plamondon
 7,870,277 B2 1/2011 Korrapati et al.
 7,916,047 B2 3/2011 Samuels et al.
 7,925,693 B2 4/2011 Swander et al.
 2001/0030970 A1 10/2001 Wiryaman et al.
 2002/0042875 A1 4/2002 Shukla
 2002/0112085 A1 8/2002 Berg
 2002/0133598 A1 9/2002 Strahm et al.
 2002/0199007 A1 12/2002 Clayton et al.
 2003/0131079 A1 7/2003 Neale et al.
 2004/0015591 A1 1/2004 Wang
 2004/0024882 A1 2/2004 Austin et al.
 2004/0177158 A1 9/2004 Bauch et al.

2004/0210320 A1 10/2004 Pandya
 2006/0080441 A1 * 4/2006 Chen et al. 709/225
 2006/0126619 A1 6/2006 Teisberg et al.
 2006/0179147 A1 8/2006 Tran et al.
 2007/0245409 A1 10/2007 Harris et al.
 2008/0034072 A1 2/2008 He et al.
 2008/0034111 A1 2/2008 Kamath et al.
 2008/0034416 A1 2/2008 Kumar et al.
 2008/0189393 A1 8/2008 Wagner
 2008/0320154 A1 12/2008 Demmer et al.
 2009/0063657 A1 3/2009 Samuels et al.
 2010/0023582 A1 1/2010 Pedersen et al.
 2010/0088398 A1 4/2010 Plamondon
 2010/0131659 A1 5/2010 Narayana et al.
 2010/0241694 A1 9/2010 Jensen et al.
 2010/0281112 A1 11/2010 Plamondon
 2010/0281162 A1 11/2010 Venkatraman et al.
 2010/0287227 A1 11/2010 Goel et al.
 2010/0318665 A1 12/2010 Demmer et al.

OTHER PUBLICATIONS

K.V.S. Mounika, Naduri Jyothirmai, and A. Rama Krishna.
 "Dynamic Routing with Security Considerations." International
 Journal of Innovative Technology and Exploring Engineering
 (IJITEE), ISSN: 2278-3075, vol. 1, Issue-6, Nov. 2012.
 Notice of Allowance for U.S. Appl. No. 11/465,950 dated Jan. 28,
 2010.
 Office Action for U.S. Appl. No. 11/462,174 dated Feb. 19, 2010.
 Office Action for U.S. Appl. No. 11/465,950 dated Jul. 2, 2009.
 Office Action for U.S. Appl. No. 12/823,643 dated Oct. 17, 2011.
 Office Action for U.S. Appl. No. 11/462,174 dated Jul. 28, 2009.
 Supplemental Notice of Allowance for U.S. Appl. No. 11/465,950
 dated Mar. 23, 2010.
 US Notice of Allowance for U.S. Appl. No. 11/462,174 dated Jun. 27,
 2013.
 US Notice of Allowance for U.S. Appl. No. 12/823,643 dated May
 18, 2012.
 US Office Action for U.S. Appl. No. 12/823,643 dated Dec. 9, 2010.
 US Office Action for U.S. Appl. No. 12/823,643 dated May 4, 2011.
 US Office Action for U.S. Appl. No. 11/462,174 dated Jan. 10, 2012.
 US Office Action for U.S. Appl. No. 11/462,174 dated Jul. 7, 2011.

* cited by examiner

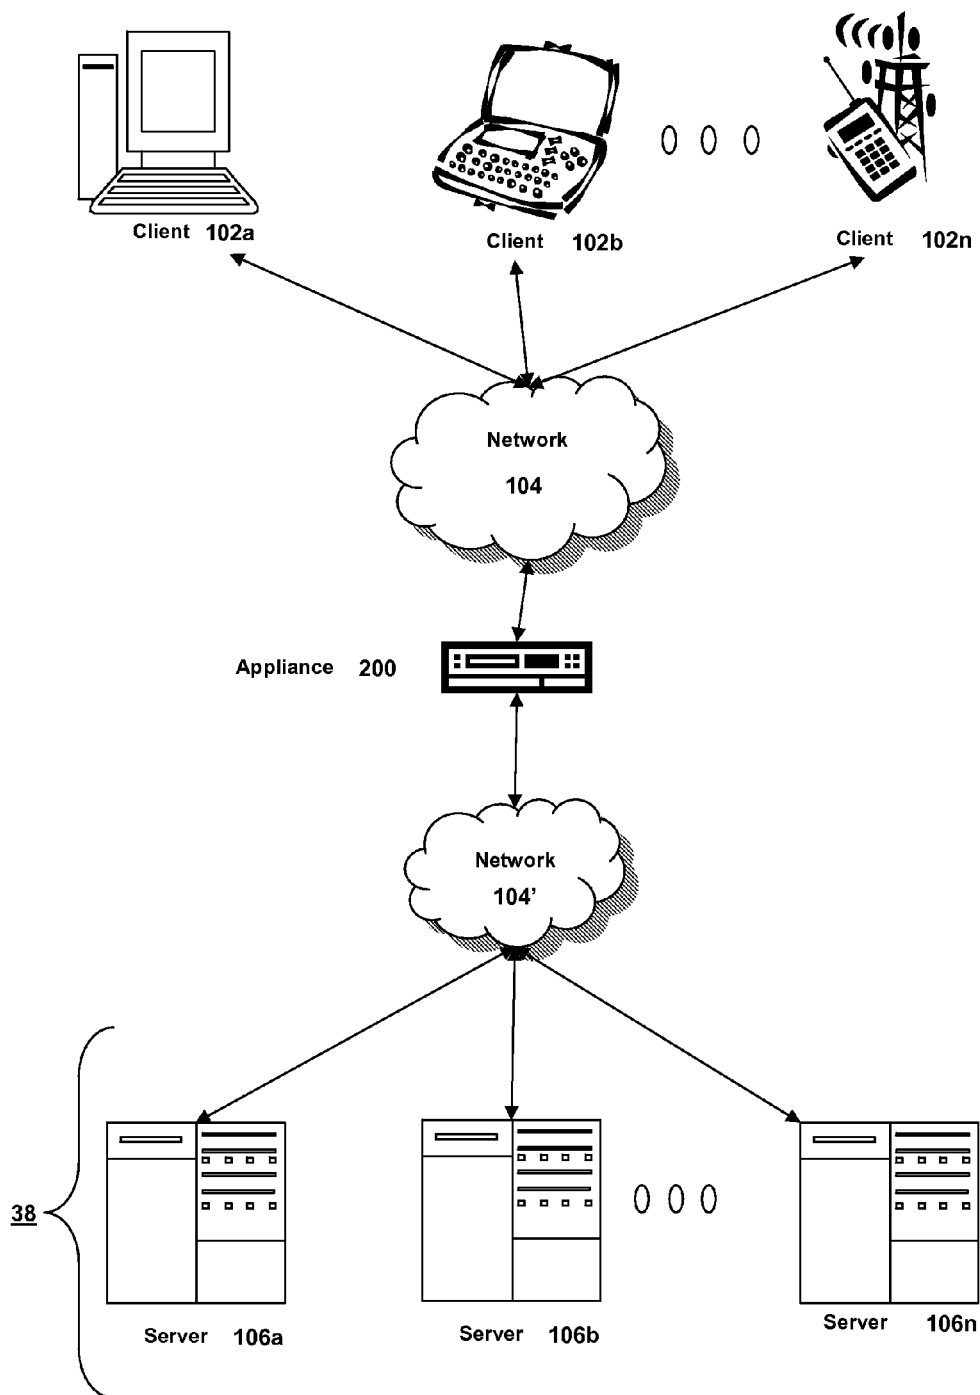

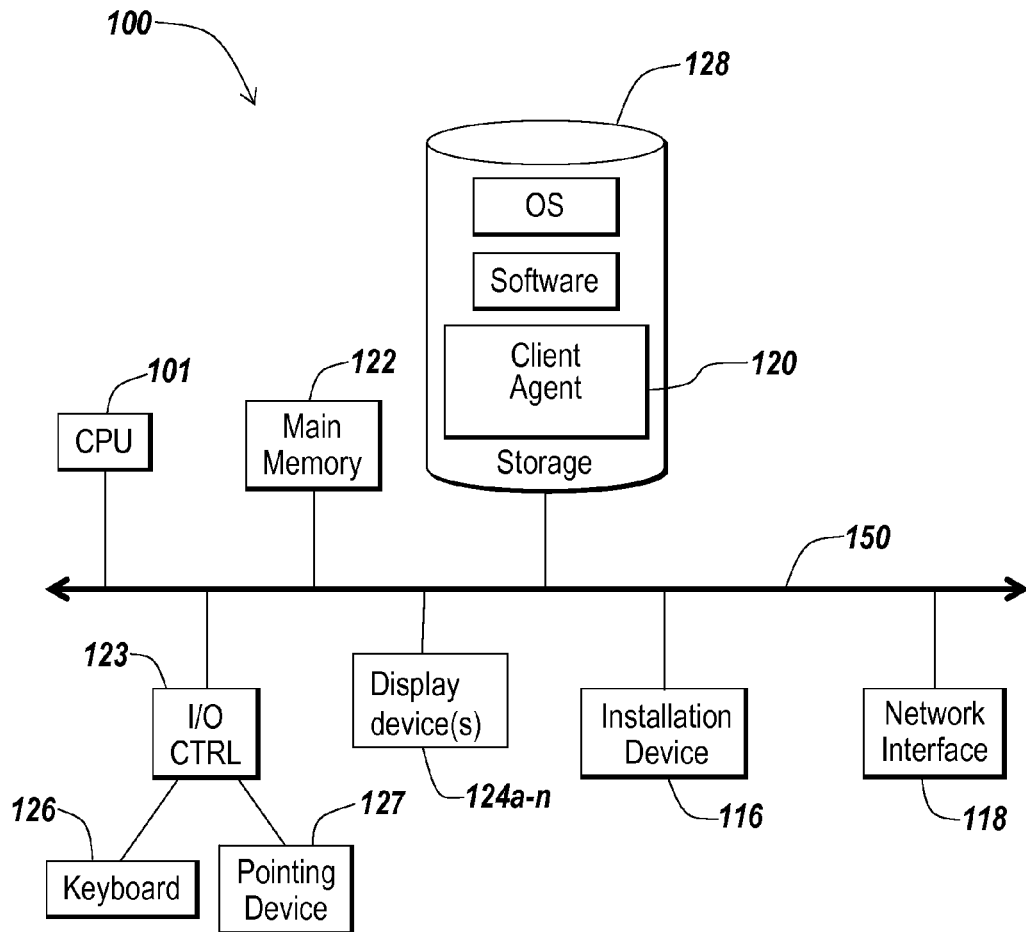

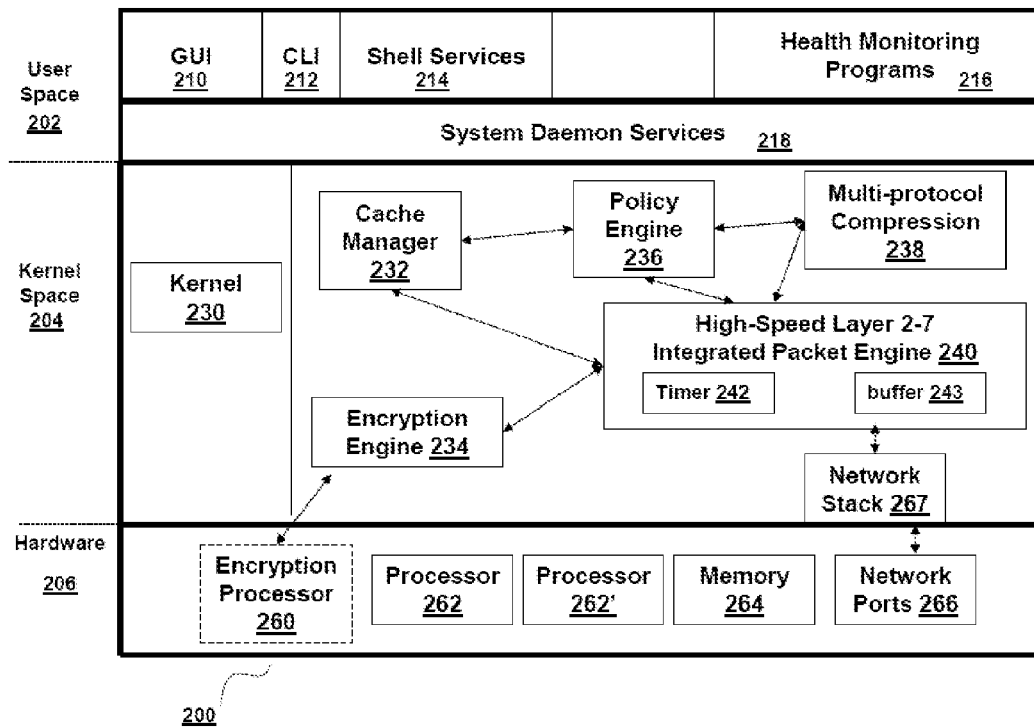



Fig. 1A


*Fig. 1B*

*Fig. 1C*

*Fig. 1D*

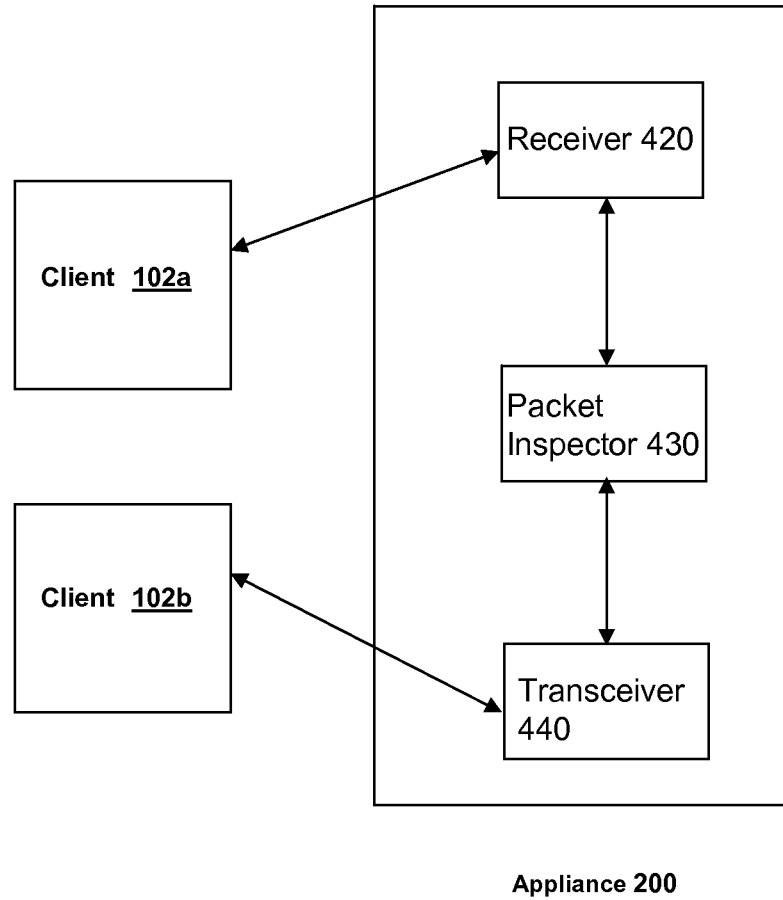
*Fig. 2A*

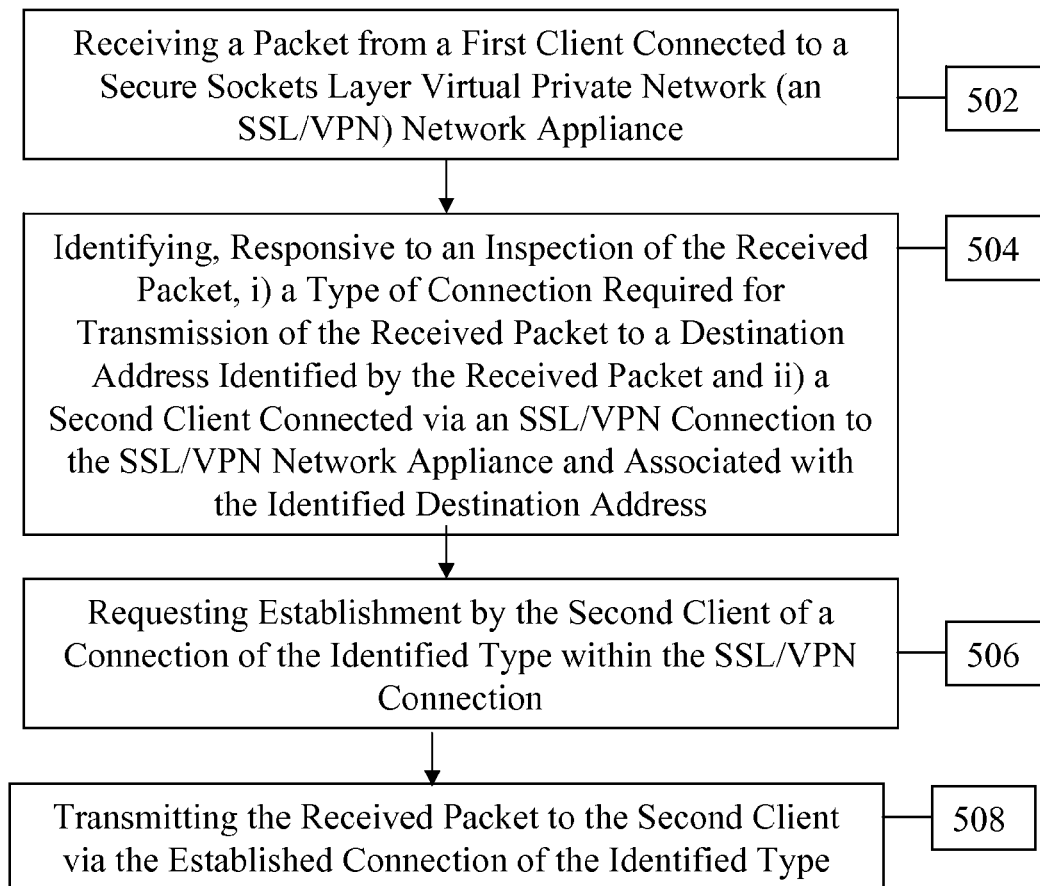
*Fig. 2B*

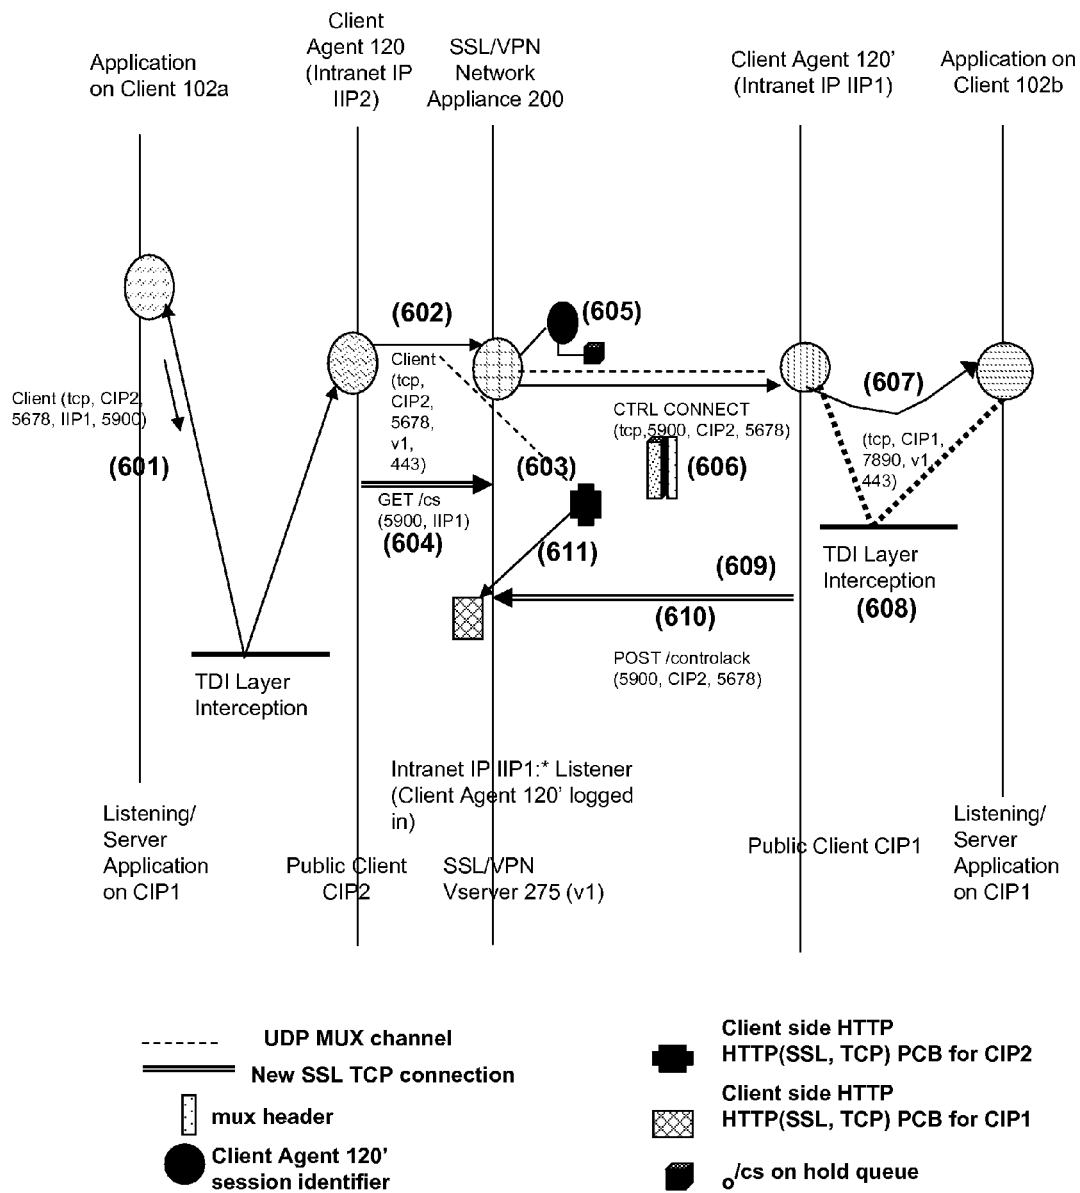
Client 102

100

Fig. 3

*Fig. 4*

*Fig. 5*

*Fig. 6*

1

METHODS AND SYSTEMS FOR ROUTING PACKETS IN A VPN-CLIENT-TO-VPN-CLIENT CONNECTION VIA AN SSL/VPN NETWORK APPLIANCE

RELATED APPLICATIONS

This U.S. patent application claims priority to and is a continuation of U.S. application Ser. No. 11/462,174, filed on Aug. 3, 2006, titled "Methods And Systems For Routing Packets In A VPN-Client-To-VPN-Client Connection Via An SSL/VPN Network Appliance," the disclosure of which is considered part of the disclosure of this application and is herein incorporated by reference in its entirety.

FIELD OF THE INVENTION

The present invention generally relates to data communication networks. In particular, the present invention relates to methods and systems for routing packets between clients on a virtual private network, via a network appliance.

BACKGROUND OF THE INVENTION

In conventional systems, virtual private networks (VPN) provide a client computer with secure access to a private network. Typically, the client is assigned an internet protocol (IP) address on the private network after a VPN server authenticates the client. Once the client is authenticated and can access the VPN, the client may securely access resources residing on the private network.

Typically, VPNs do not support peer-to-peer connections between clients on the VPN. Some systems enable a client to establish a separate, secure connection to other clients when the client wishes to transmit data to another client, effectively creating a second virtual private network. In some of these systems, the clients communicate directly with each other through a connection that is independent of the virtual private network. But conventional systems do not typically enable clients who have already established connections to the VPN to communicate with each other in a peer-to-peer mode.

A system in which a VPN server enabled clients connected to a virtual private network to communicate with each other in a peer-to-peer mode while utilizing existing connections to the VPN would be desirable.

BRIEF SUMMARY OF THE INVENTION

In one aspect of the invention, a method for routing packets between clients includes the step of receiving a packet from a first client connected to a secure sockets layer virtual private network (an SSL/VPN) network appliance. An identification is made, responsive to an inspection of the received packet, of i) a type of connection required for transmission of the received packet to a destination address identified by the received packet and ii) a second client connected via an SSL/VPN connection to the SSL/VPN network appliance and associated with the identified destination address. A request is made for establishment by the second client of a connection of the identified type within the SSL/VPN connection. The received packet is transmitted to the second client via the established connection of the identified type.

In one embodiment, the step of receiving the packet includes receiving, by the SSL/VPN network appliance, the packet via a secure sockets layer virtual private network (an SSL VPN) connection. In another embodiment, the step of

2

identifying the type of connection and the second client includes inspecting, by the SSL/VPN network appliance, the received packet. In still another embodiment, the step of identifying the type of connection includes identifying, responsive to the inspection of the received packet, a type of connection comprising a TCP connection. In yet another embodiment, the step of identifying the type of connection includes identifying, responsive to the inspection of the received packet, a type of connection comprising a UDP connection.

In one embodiment, the SSL/VPN network appliance requests establishment by the second client of a connection of the identified type. In another embodiment, the connection requested is a connection of the identified type between the second client and the SSL/VPN network appliance. In still another embodiment, the second client establishes a connection of the identified type.

In another aspect, a system for routing packets between clients comprises a secure sockets layer virtual private network (an SSL/VPN) network appliance comprising a receiver, a packet inspector, and a transceiver. The receiver receives a packet from a first client over a connection. The packet inspector identifies, responsive to an inspection of the received packet, 1) a type of connection required for transmission of the received packet to a destination address identified by the received packet and 2) a second client connected via an SSL/VPN connection to the SSL/VPN network appliance and associated with the identified destination address. The transceiver requests establishment of a connection of the identified type and transmits the received packet to the second client over the requested connection.

In still another aspect, a system for routing packets between clients comprises a first client transmitting a packet, a second client, and a secure sockets layer virtual private network (an SSL/VPN) network appliance. The SSL/VPN network appliance comprises a receiver, a packet inspector, and a transceiver. The receiver receives the packet from the first client over a connection. The packet inspector identifies, responsive to an inspection of the received packet, 1) a type of connection required for transmission of the received packet to a destination address identified by the received packet and 2) the second client. The transceiver requests establishment by the second client of a connection of the identified type and transmits the received packet to the second client over the requested connection.

In one embodiment, the first client establishes a secure sockets layer virtual private network (an SSL/VPN) connection to the SSL/VPN network appliance. In another embodiment, the first client transmits the packet to the SSL/VPN network appliance over a secure sockets layer virtual private network (an SSL/VPN) connection.

In one embodiment, the second client further comprises a connection via an SSL/VPN connection to the SSL/VPN network appliance. In another embodiment, the second client further comprises an association with the destination address identified by the received packet. In still another embodiment, the second client establishes a connection of the identified type between the second client and the SSL/VPN network appliance. In yet another embodiment, the second client establishes a connection of the identified type within an SSL/VPN connection to the SSL/VPN network appliance.

The details of various embodiments of the invention are set forth in the accompanying drawings and the description below.

BRIEF DESCRIPTION OF THE FIGURES

The foregoing and other objects, aspects, features, and advantages of the invention will become more apparent and

3

better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

FIG. 1A is a block diagram of an embodiment of a network environment for a client to access a server via an appliance;

FIG. 1B is a block diagram of an embodiment of an environment for delivering a computing environment from a server to a client via an appliance;

FIGS. 1C and 1D are block diagrams of embodiments of a computing device;

FIG. 2A is a block diagram of an embodiment of an appliance for processing communications between a client and a server;

FIG. 2B is a block diagram of another embodiment of an appliance for optimizing, accelerating, load-balancing and routing communications between a client and a server;

FIG. 3 is a block diagram of an embodiment of a client for communicating with a server via the appliance;

FIG. 4 is a block diagram depicting one embodiment of a system for routing packets between clients;

FIG. 5 is a flow diagram depicting one embodiment of the steps taken in a method for routing packets between clients; and

FIG. 6 is a flow diagram depicting another embodiment of the steps taken in a method for routing packets between clients.

The features and advantages of the present invention will become more apparent from the detailed description set forth below when taken in conjunction with the drawings, in which like reference characters identify corresponding elements throughout. In the drawings, like reference numbers generally indicate identical, functionally similar, and/or structurally similar elements.

DETAILED DESCRIPTION OF THE INVENTION

Prior to discussing the specifics of embodiments of the systems and methods of an appliance and/or client, it may be helpful to discuss the network and computing environments in which such embodiments may be deployed. Referring now to FIG. 1A, an embodiment of a network environment is depicted. In brief overview, the network environment comprises one or more clients **102a-102n** (also generally referred to as local machine(s) **102**, or client(s) **102**) in communication with one or more servers **106a-106n** (also generally referred to as server(s) **106**, or remote machine(s) **106**) via one or more networks **104, 104'** (generally referred to as network **104**). In some embodiments, a client **102** communicates with a server **106** via an appliance **200**.

Although FIG. 1A shows a network **104** and a network **104'** between the clients **102** and the servers **106**, the clients **102** and the servers **106** may be on the same network **104**. The networks **104** and **104'** can be the same type of network or different types of networks. The network **104** and/or the network **104'** can be a local-area network (LAN), such as a company Intranet, a metropolitan area network (MAN), or a wide area network (WAN), such as the Internet or the World Wide Web. In one embodiment, network **104'** may be a private network and network **104** may be a public network. In some embodiments, network **104** may be a private network and network **104'** a public network. In another embodiment, networks **104** and **104'** may both be private networks. In some embodiments, clients **102** may be located at a branch office of a corporate enterprise communicating via a WAN connection over the network **104** to the servers **106** located at a corporate data center.

4

The network **104** and/or **104'** be any type and/or form of network and may include any of the following: a point to point network, a broadcast network, a wide area network, a local area network, a telecommunications network, a data communication network, a computer network, an ATM (Asynchronous Transfer Mode) network, a SONET (Synchronous Optical Network) network, a SDH (Synchronous Digital Hierarchy) network, a wireless network and a wireline network. In some embodiments, the network **104** may comprise a wireless link, such as an infrared channel or satellite band. The topology of the network **104** and/or **104'** may be a bus, star, or ring network topology. The network **104** and/or **104'** and network topology may be of any such network or network topology as known to those ordinarily skilled in the art capable of supporting the operations described herein.

As shown in FIG. 1A, the appliance **200**, which also may be referred to as an interface unit **200** or gateway **200**, is shown between the networks **104** and **104'**. In some embodiments, the appliance **200** may be located on network **104**. For example, a branch office of a corporate enterprise may deploy an appliance **200** at the branch office. In other embodiments, the appliance **200** may be located on network **104'**. For example, an appliance **200** may be located at a corporate data center. In yet another embodiment, a plurality of appliances **200** may be deployed on network **104**. In some embodiments, a plurality of appliances **200** may be deployed on network **104'**. In one embodiment, a first appliance **200** communicates with a second appliance **200'**. In other embodiments, the appliance **200** could be a part of any client **102** or server **106** on the same or different network **104, 104'** as the client **102**. One or more appliances **200** may be located at any point in the network or network communications path between a client **102** and a server **106**.

In one embodiment, the system may include multiple, logically-grouped servers **106**. In these embodiments, the logical group of servers may be referred to as a server farm **38**. In some of these embodiments, the servers **106** may be geographically dispersed. In some cases, a farm **38** may be administered as a single entity. In other embodiments, the server farm **38** comprises a plurality of server farms **38**. In one embodiment, the server farm executes one or more applications on behalf of one or more clients **102**.

The servers **106** within each farm **38** can be heterogeneous. One or more of the servers **106** can operate according to one type of operating system platform (e.g., WINDOWS NT, manufactured by Microsoft Corp. of Redmond, Wash.), while one or more of the other servers **106** can operate on according to another type of operating system platform (e.g., Unix or Linux). The servers **106** of each farm **38** do not need to be physically proximate to another server **106** in the same farm **38**. Thus, the group of servers **106** logically grouped as a farm **38** may be interconnected using a wide-area network (WAN) connection or medium-area network (MAN) connection. For example, a farm **38** may include servers **106** physically located in different continents or different regions of a continent, country, state, city, campus, or room. Data transmission speeds between servers **106** in the farm **38** can be increased if the servers **106** are connected using a local-area network (LAN) connection or some form of direct connection.

Servers **106** may be referred to as a file server, application server, web server, proxy server, or gateway server. In some embodiments, a server **106** may have the capacity to function as either an application server or as a master application server. In one embodiment, a server **106** may include an Active Directory. The clients **102** may also be referred to as client nodes or endpoints. In some embodiments, a client **102** has the capacity to function as both a client node seeking

5

access to applications on a server and as an application server providing access to hosted applications for other clients **102a-102n**.

In some embodiments, a client **102** communicates with a server **106**. In one embodiment, the client **102** communicates directly with one of the servers **106** in a farm **38**. In another embodiment, the client **102** executes a program neighborhood application to communicate with a server **106** in a farm **38**. In still another embodiment, the server **106** provides the functionality of a master node. In some embodiments, the client **102** communicates with the server **106** in the farm **38** through a network **104**. Over the network **104**, the client **102** can, for example, request execution of various applications hosted by the servers **106a-106n** in the farm **38** and receive output of the results of the application execution for display. In some embodiments, only the master node provides the functionality required to identify and provide address information associated with a server **106'** hosting a requested application.

In one embodiment, the server **106** provides functionality of a web server. In another embodiment, the server **106a** receives requests from the client **102**, forwards the requests to a second server **106b** and responds to the request by the client **102** with a response to the request from the server **106b**. In still another embodiment, the server **106** acquires an enumeration of applications available to the client **102** and address information associated with a server **106** hosting an application identified by the enumeration of applications. In yet another embodiment, the server **106** presents the response to the request to the client **102** using a web interface. In one embodiment, the client **102** communicates directly with the server **106** to access the identified application. In another embodiment, the client **102** receives application output data, such as display data, generated by an execution of the identified application on the server **106**.

Referring now to FIG. 1B, a network environment for delivering and/or operating a computing environment on a client **102** is depicted. In some embodiments, a server **106** includes an application delivery system **190** for delivering a computing environment or an application and/or data file to one or more clients **102**. In brief overview, a client **10** is in communication with a server **106** via network **104**, **104'** and appliance **200**. For example, the client **102** may reside in a remote office of a company, e.g., a branch office, and the server **106** may reside at a corporate data center. The client **102** comprises a client agent **120**, and a computing environment **15**. The computing environment **15** may execute or operate an application that accesses, processes or uses a data file. The computing environment **15**, application and/or data file may be delivered via the appliance **200** and/or the server **106**.

In some embodiments, the appliance **200** accelerates delivery of a computing environment **15**, or any portion thereof, to a client **102**. In one embodiment, the appliance **200** accelerates the delivery of the computing environment **15** by the application delivery system **190**. For example, the embodiments described herein may be used to accelerate delivery of a streaming application and data file processable by the application from a central corporate data center to a remote user location, such as a branch office of the company. In another embodiment, the appliance **200** accelerates transport layer traffic between a client **102** and a server **106**. The appliance **200** may provide acceleration techniques for accelerating any transport layer payload from a server **106** to a client **102**, such as: 1) transport layer connection pooling, 2) transport layer connection multiplexing, 3) transport control protocol buffering, 4) compression and 5) caching. In some embodiments,

6

the appliance **200** provides load balancing of servers **106** in responding to requests from clients **102**. In other embodiments, the appliance **200** acts as a proxy or access server to provide access to the one or more servers **106**. In another embodiment, the appliance **200** provides a secure virtual private network connection from a first network **104** of the client **102** to the second network **104'** of the server **106**, such as an SSL VPN connection. In yet other embodiments, the appliance **200** provides application firewall security, control and management of the connection and communications between a client **102** and a server **106**.

In some embodiments, the application delivery management system **190** provides application delivery techniques to deliver a computing environment to a desktop of a user, remote or otherwise, based on a plurality of execution methods and based on any authentication and authorization policies applied via a policy engine **195**. With these techniques, a remote user may obtain a computing environment and access to server stored applications and data files from any network connected device **100**. In one embodiment, the application delivery system **190** may reside or execute on a server **106**. In another embodiment, the application delivery system **190** may reside or execute on a plurality of servers **106a-106n**. In some embodiments, the application delivery system **190** may execute in a server farm **38**. In one embodiment, the server **106** executing the application delivery system **190** may also store or provide the application and data file. In another embodiment, a first set of one or more servers **106** may execute the application delivery system **190**, and a different server **106n** may store or provide the application and data file. In some embodiments, each of the application delivery system **190**, the application, and data file may reside or be located on different servers. In yet another embodiment, any portion of the application delivery system **190** may reside, execute or be stored on or distributed to the appliance **200**, or a plurality of appliances.

The client **102** may include a computing environment **15** for executing an application that uses or processes a data file. The client **102** via networks **104**, **104'** and appliance **200** may request an application and data file from the server **106**. In one embodiment, the appliance **200** may forward a request from the client **102** to the server **106**. For example, the client **102** may not have the application and data file stored or accessible locally. In response to the request, the application delivery system **190** and/or server **106** may deliver the application and data file to the client **102**. For example, in one embodiment, the server **106** may transmit the application as an application stream to operate in computing environment **15** on client **102**.

In some embodiments, the application delivery system **190** comprises any portion of the Citrix Access Suite™ by Citrix Systems, Inc., such as the MetaFrame or Citrix Presentation Server™ and/or any of the Microsoft® Windows Terminal Services manufactured by the Microsoft Corporation. In one embodiment, the application delivery system **190** may deliver one or more applications to clients **102** or users via a remote-display protocol or otherwise via remote-based or server-based computing. In another embodiment, the application delivery system **190** may deliver one or more applications to clients or users via streaming of the application.

In one embodiment, the application delivery system **190** includes a policy engine **195** for controlling and managing the access to, selection of application execution methods and the delivery of applications. In some embodiments, the policy engine **195** determines the one or more applications a user or client **102** may access. In another embodiment, the policy engine **195** determines how the application should be delivered to the user or client **102**, e.g., the method of execution. In

7

some embodiments, the application delivery system **190** provides a plurality of delivery techniques from which to select a method of application execution, such as a server-based computing, streaming or delivering the application locally to the client **120** for local execution.

In one embodiment, a client **102** requests execution of an application program and the application delivery system **190** comprising a server **106** selects a method of executing the application program. In some embodiments, the server **106** receives credentials from the client **102**. In another embodiment, the server **106** receives a request for an enumeration of available applications from the client **102**. In one embodiment, in response to the request or receipt of credentials, the application delivery system **190** enumerates a plurality of application programs available to the client **102**. The application delivery system **190** receives a request to execute an enumerated application. The application delivery system **190** selects one of a predetermined number of methods for executing the enumerated application, for example, responsive to a policy of a policy engine. The application delivery system **190** may select a method of execution of the application enabling the client **102** to receive application-output data generated by execution of the application program on a server **106**. The application delivery system **190** may select a method of execution of the application enabling the local machine **10** to execute the application program locally after retrieving a plurality of application files comprising the application. In yet another embodiment, the application delivery system **190** may select a method of execution of the application to stream the application via the network **104** to the client **102**.

A client **102** may execute, operate or otherwise provide an application, which can be any type and/or form of software, program, or executable instructions such as any type and/or form of web browser, web-based client, client-server application, a thin-client computing client, an ActiveX control, or a Java applet, or any other type and/or form of executable instructions capable of executing on client **102**. In some embodiments, the application may be a server-based or a remote-based application executed on behalf of the client **102** on a server **106**. In one embodiment the server **106** may display output to the client **102** using any thin-client or remote-display protocol, such as the Independent Computing Architecture (ICA) protocol manufactured by Citrix Systems, Inc. of Ft. Lauderdale, Fla. or the Remote Desktop Protocol (RDP) manufactured by the Microsoft Corporation of Redmond, Wash. The application can use any type of protocol and it can be, for example, an HTTP client, an FTP client, an Oscar client, or a Telnet client. In other embodiments, the application comprises any type of software related to VoIP communications, such as a soft IP telephone. In further embodiments, the application comprises any application related to real-time data communications, such as applications for streaming video and/or audio.

In some embodiments, the server **106** or a server farm **38** may be running one or more applications, such as an application providing a thin-client computing or remote display presentation application. In one embodiment, the server **106** or server farm **38** executes as an application, any portion of the Citrix Access Suite™ by Citrix Systems, Inc., such as the MetaFrame or Citrix Presentation Server™, and/or any of the Microsoft® Windows Terminal Services manufactured by the Microsoft Corporation. In one embodiment, the application is an ICA client, developed by Citrix Systems, Inc. of Fort Lauderdale, Fla. In other embodiments, the application includes a Remote Desktop (RDP) client, developed by Microsoft Corporation of Redmond, Wash. Also, the server

8

106 may run an application, which for example, may be an application server providing email services such as Microsoft Exchange manufactured by the Microsoft Corporation of Redmond, Wash., a web or Internet server, or a desktop sharing server, or a collaboration server. In some embodiments, any of the applications may comprise any type of hosted service or products, such as GoToMeeting™ provided by Citrix Online Division, Inc. of Santa Barbara, Calif., WebEx™ provided by WebEx, Inc. of Santa Clara, Calif., or Microsoft Office Live Meeting provided by Microsoft Corporation of Redmond, Wash.

The client **102**, server **106**, and appliance **200** may be deployed as and/or executed on any type and form of computing device, such as a computer, network device or appliance capable of communicating on any type and form of network and performing the operations described herein. FIGS. 1C and 1D depict block diagrams of a computing device **100** useful for practicing an embodiment of the client **102**, server **106** or appliance **200**. As shown in FIGS. 1C and 1D, each computing device **100** includes a central processing unit **101**, and a main memory unit **122**. As shown in FIG. 1C, a computing device **100** may include a visual display device **124**, a keyboard **126** and/or a pointing device **127**, such as a mouse. Each computing device **100** may also include additional optional elements, such as one or more input/output devices **130a-130b** (generally referred to using reference numeral **130**), and a cache memory **140** in communication with the central processing unit **101**.

The central processing unit **101** is any logic circuitry that responds to and processes instructions fetched from the main memory unit **122**. In many embodiments, the central processing unit is provided by a microprocessor unit, such as: those manufactured by Intel Corporation of Mountain View, Calif.; those manufactured by Motorola Corporation of Schaumburg, Ill.; those manufactured by Transmeta Corporation of Santa Clara, Calif.; the RS/6000 processor, those manufactured by International Business Machines of White Plains, N.Y.; or those manufactured by Advanced Micro Devices of Sunnyvale, Calif. The computing device **100** may be based on any of these processors, or any other processor capable of operating as described herein.

Main memory unit **122** may be one or more memory chips capable of storing data and allowing any storage location to be directly accessed by the microprocessor **101**, such as Static random access memory (SRAM), Burst SRAM or Synchronous Burst SRAM (BSRAM), Dynamic random access memory (DRAM), Fast Page Mode DRAM (FPM DRAM), Enhanced DRAM (EDRAM), Extended Data Output RAM (EDO RAM), Extended Data Output DRAM (EDO DRAM), Burst Extended Data Output DRAM (BEDO DRAM), Enhanced DRAM (EDRAM), synchronous DRAM (SDRAM), JEDEC SDRAM, PC100 SDRAM, Double Data Rate SDRAM (DDR SDRAM), Enhanced SDRAM (ESDRAM), SyncLink DRAM (SLDRAM), Direct Rambus DRAM (DRDRAM), or Ferroelectric RAM (FRAM). The main memory **122** may be based on any of the above described memory chips, or any other available memory chips capable of operating as described herein. In the embodiment shown in FIG. 1C, the processor **101** communicates with main memory **122** via a system bus **150** (described in more detail below). FIG. 1C depicts an embodiment of a computing device **100** in which the processor communicates directly with main memory **122** via a memory port **103**. For example, in FIG. 1D the main memory **122** may be DRDRAM.

FIG. 1D depicts an embodiment in which the main processor **101** communicates directly with cache memory **140** via a secondary bus, sometimes referred to as a backside bus. In

other embodiments, the main processor **101** communicates with cache memory **140** using the system bus **150**. Cache memory **140** typically has a faster response time than main memory **122** and is typically provided by SRAM, BSRAM, or EDRAM. In the embodiment shown in FIG. 1C, the processor **101** communicates with various I/O devices **130** via a local system bus **150**. Various busses may be used to connect the central processing unit **101** to any of the I/O devices **130**, including a VESA VL bus, an ISA bus, an EISA bus, a MicroChannel Architecture (MCA) bus, a PCI bus, a PCI-X bus, a PCI-Express bus, or a NuBus. For embodiments in which the I/O device is a video display **124**, the processor **101** may use an Advanced Graphics Port (AGP) to communicate with the display **124**. FIG. 1D depicts an embodiment of a computer **100** in which the main processor **101** communicates directly with I/O device **130** via HyperTransport, Rapid I/O, or InfiniBand. FIG. 1D also depicts an embodiment in which local busses and direct communication are mixed: the processor **101** communicates with I/O device **130** using a local interconnect bus while communicating with I/O device **130** directly.

The computing device **100** may support any suitable installation device **116**, such as a floppy disk drive for receiving floppy disks such as 3.5-inch, 5.25-inch disks or ZIP disks, a CD-ROM drive, a CD-R/RW drive, a DVD-ROM drive, tape drives of various formats, USB device, hard-drive or any other device suitable for installing software and programs such as any client agent **120**, or portion thereof. The computing device **100** may further comprise a storage device **128**, such as one or more hard disk drives or redundant arrays of independent disks, for storing an operating system and other related software, and for storing application software programs such as any program related to the client agent **120**. Optionally, any of the installation devices **116** could also be used as the storage device **128**. Additionally, the operating system and the software can be run from a bootable medium, for example, a bootable CD, such as KNOPPIX®, a bootable CD for GNU/Linux that is available as a GNU/Linux distribution from knoppix.net.

Furthermore, the computing device **100** may include a network interface **118** to interface to a Local Area Network (LAN), Wide Area Network (WAN) or the Internet through a variety of connections including, but not limited to, standard telephone lines, LAN or WAN links (e.g., 802.11, T1, T3, 56 kb, X.25), broadband connections (e.g., ISDN, Frame Relay, ATM), wireless connections, or some combination of any or all of the above. The network interface **118** may comprise a built-in network adapter, network interface card, PCMCIA network card, card bus network adapter, wireless network adapter, USB network adapter, modem or any other device suitable for interfacing the computing device **100** to any type of network capable of communication and performing the operations described herein.

A wide variety of I/O devices **130a-130n** may be present in the computing device **100**. Input devices include keyboards, mice, trackpads, trackballs, microphones, and drawing tablets. Output devices include video displays, speakers, inkjet printers, laser printers, and dye-sublimation printers. The I/O devices **130** may be controlled by an I/O controller **123** as shown in FIG. 1C. The I/O controller may control one or more I/O devices such as a keyboard **126** and a pointing device **127**, e.g., a mouse or optical pen. Furthermore, an I/O device may also provide storage **128** and/or an installation medium **116** for the computing device **100**. In still other embodiments, the computing device **100** may provide USB connections to

receive handheld USB storage devices such as the USB Flash Drive line of devices manufactured by Twintech Industry, Inc. of Los Alamitos, Calif.

In some embodiments, the computing device **100** may comprise or be connected to multiple display devices **124a-124n**, which each may be of the same or different type and/or form. As such, any of the I/O devices **130a-130n** and/or the I/O controller **123** may comprise any type and/or form of suitable hardware, software, or combination of hardware and software to support, enable or provide for the connection and use of multiple display devices **124a-124n** by the computing device **100**. For example, the computing device **100** may include any type and/or form of video adapter, video card, driver, and/or library to interface, communicate, connect or otherwise use the display devices **124a-124n**. In one embodiment, a video adapter may comprise multiple connectors to interface to multiple display devices **124a-124n**. In other embodiments, the computing device **100** may include multiple video adapters, with each video adapter connected to one or more of the display devices **124a-124n**. In some embodiments, any portion of the operating system of the computing device **100** may be configured for using multiple displays **124a-124n**. In other embodiments, one or more of the display devices **124a-124n** may be provided by one or more other computing devices, such as computing devices **100a** and **100b** connected to the computing device **100**, for example, via a network. These embodiments may include any type of software designed and constructed to use another computer's display device as a second display device **124a** for the computing device **100**. One ordinarily skilled in the art will recognize and appreciate the various ways and embodiments that a computing device **100** may be configured to have multiple display devices **124a-124n**.

In further embodiments, an I/O device **130** may be a bridge **170** between the system bus **150** and an external communication bus, such as a USB bus, an Apple Desktop Bus, an RS-232 serial connection, a SCSI bus, a FireWire bus, a FireWire 800 bus, an Ethernet bus, an AppleTalk bus, a Gigabit Ethernet bus, an Asynchronous Transfer Mode bus, a HIPPI bus, a Super HIPPI bus, a SerialPlus bus, a SCSI/LAMP bus, a FibreChannel bus, or a Serial Attached small computer system interface bus.

A computing device **100** of the sort depicted in FIGS. 1C and 1D typically operate under the control of operating systems, which control scheduling of tasks and access to system resources. The computing device **100** can be running any operating system such as any of the versions of the Microsoft® Windows operating systems, the different releases of the Unix and Linux operating systems, any version of the Mac OS® for Macintosh computers, any embedded operating system, any real-time operating system, any open source operating system, any proprietary operating system, any operating systems for mobile computing devices, or any other operating system capable of running on the computing device and performing the operations described herein. Typical operating systems include: WINDOWS 3.x, WINDOWS 95, WINDOWS 98, WINDOWS 2000, WINDOWS NT 3.51, WINDOWS NT 4.0, WINDOWS CE, and WINDOWS XP, all of which are manufactured by Microsoft Corporation of Redmond, Wash.; MacOS, manufactured by Apple Computer of Cupertino, Calif.; OS/2, manufactured by International Business Machines of Armonk, N.Y.; and Linux, a freely-available operating system distributed by Caldera Corp. of Salt Lake City, Utah, or any type and/or form of a Unix operating system, among others.

In other embodiments, the computing device **100** may have different processors, operating systems, and input devices

11

consistent with the device. For example, in one embodiment the computer **100** is a Treo 180, 270, 1060, 600 or 650 smart phone manufactured by Palm, Inc. In this embodiment, the Treo smart phone is operated under the control of the PalmOS operating system and includes a stylus input device as well as a five-way navigator device. Moreover, the computing device **100** can be any workstation, desktop computer, laptop or notebook computer, server, handheld computer, mobile telephone, any other computer, or other form of computing or telecommunications device that is capable of communication and that has sufficient processor power and memory capacity to perform the operations described herein.

FIG. 2A illustrates an example embodiment of the appliance **200**. The architecture of the appliance **200** in FIG. 2A is provided by way of illustration only and is not intended to be limiting. As shown in FIG. 2, appliance **200** comprises a hardware layer **206** and a software layer divided into a user space **202** and a kernel space **204**.

Hardware layer **206** provides the hardware elements upon which programs and services within kernel space **204** and user space **202** are executed. Hardware layer **206** also provides the structures and elements which allow programs and services within kernel space **204** and user space **202** to communicate data both internally and externally with respect to appliance **200**. As shown in FIG. 2, the hardware layer **206** includes a processing unit **262** for executing software programs and services, a memory **264** for storing software and data, network ports **266** for transmitting and receiving data over a network, and an encryption processor **260** for performing functions related to Secure Sockets Layer processing of data transmitted and received over the network. In some embodiments, the central processing unit **262** may perform the functions of the encryption processor **260** in a single processor. Additionally, the hardware layer **206** may comprise multiple processors for each of the processing unit **262** and the encryption processor **260**. The processor **262** may include any of the processors **101** described above in connection with FIGS. 1C and 1D. In some embodiments, the central processing unit **262** may perform the functions of the encryption processor **260** in a single processor. Additionally, the hardware layer **206** may comprise multiple processors for each of the processing unit **262** and the encryption processor **260**. For example, in one embodiment, the appliance **200** comprises a first processor **262** and a second processor **262'**. In other embodiments, the processor **262** or **262'** comprises a multi-core processor.

Although the hardware layer **206** of appliance **200** is generally illustrated with an encryption processor **260**, processor **260** may be a processor for performing functions related to any encryption protocol, such as the Secure Socket Layer (SSL) or Transport Layer Security (TLS) protocol. In some embodiments, the processor **260** may be a general purpose processor (GPP), and in further embodiments, may be have executable instructions for performing processing of any security related protocol.

Although the hardware layer **206** of appliance **200** is illustrated with certain elements in FIG. 2, the hardware portions or components of appliance **200** may comprise any type and form of elements, hardware or software, of a computing device, such as the computing device **100** illustrated and discussed herein in conjunction with FIGS. 1C and 1D. In some embodiments, the appliance **200** may comprise a server, gateway, router, switch, bridge or other type of computing or network device, and have any hardware and/or software elements associated therewith.

The operating system of appliance **200** allocates, manages, or otherwise segregates the available system memory into

12

kernel space **204** and user space **204**. In example software architecture **200**, the operating system may be any type and/or form of Unix operating system although the invention is not so limited. As such, the appliance **200** can be running any operating system such as any of the versions of the Microsoft® Windows operating systems, the different releases of the Unix and Linux operating systems, any version of the Mac OS® for Macintosh computers, any embedded operating system, any network operating system, any real-time operating system, any open source operating system, any proprietary operating system, any operating systems for mobile computing devices or network devices, or any other operating system capable of running on the appliance **200** and performing the operations described herein.

The kernel space **204** is reserved for running the kernel **230**, including any device drivers, kernel extensions or other kernel related software. As known to those skilled in the art, the kernel **230** is the core of the operating system, and provides access, control, and management of resources and hardware-related elements of the application **104**. In accordance with an embodiment of the appliance **200**, the kernel space **204** also includes a number of network services or processes working in conjunction with a cache manager **232**, sometimes also referred to as the integrated cache, the benefits of which are described in detail further herein. Additionally, the embodiment of the kernel **230** will depend on the embodiment of the operating system installed, configured, or otherwise used by the device **200**.

In one embodiment, the device **200** comprises one network stack **267**, such as a TCP/IP based stack, for communicating with the client **102** and/or the server **106**. In one embodiment, the network stack **267** is used to communicate with a first network, such as network **108**, and a second network **110**. In some embodiments, the device **200** terminates a first transport layer connection, such as a TCP connection of a client **102**, and establishes a second transport layer connection to a server **106** for use by the client **102**, e.g., the second transport layer connection is terminated at the appliance **200** and the server **106**. The first and second transport layer connections may be established via a single network stack **267**. In other embodiments, the device **200** may comprise multiple network stacks, for example **267** and **267'**, and the first transport layer connection may be established or terminated at one network stack **267**, and the second transport layer connection on the second network stack **267'**. For example, one network stack may be for receiving and transmitting network packet on a first network, and another network stack for receiving and transmitting network packets on a second network. In one embodiment, the network stack **267** comprises a buffer **243** for queuing one or more network packets for transmission by the appliance **200**.

As shown in FIG. 2, the kernel space **204** includes the cache manager **232**, a high-speed layer 2-7 integrated packet engine **240**, an encryption engine **234**, a policy engine **236** and multi-protocol compression logic **238**. Running these components or processes **232**, **240**, **234**, **236** and **238** in kernel space **204** or kernel mode instead of the user space **202** improves the performance of each of these components, alone and in combination. Kernel operation means that these components or processes **232**, **240**, **234**, **236** and **238** run in the core address space of the operating system of the device **200**. For example, running the encryption engine **234** in kernel mode improves encryption performance by moving encryption and decryption operations to the kernel, thereby reducing the number of transitions between the memory space or a kernel thread in kernel mode and the memory space or a thread in user mode. For example, data obtained in kernel mode may not need to be

13

passed or copied to a process or thread running in user mode, such as from a kernel level data structure to a user level data structure. In another aspect, the number of context switches between kernel mode and user mode are also reduced. Additionally, synchronization of and communications between any of the components or processes **232**, **240**, **235**, **236** and **238** can be performed more efficiently in the kernel space **204**.

In some embodiments, any portion of the components **232**, **240**, **234**, **236** and **238** may run or operate in the kernel space **204**, while other portions of these components **232**, **240**, **234**, **236** and **238** may run or operate in user space **202**. In one embodiment, the appliance **200** uses a kernel-level data structure providing access to any portion of one or more network packets, for example, a network packet comprising a request from a client **102** or a response from a server **106**. In some embodiments, the kernel-level data structure may be obtained by the packet engine **240** via a transport layer driver interface or filter to the network stack **267**. The kernel-level data structure may comprise any interface and/or data accessible via the kernel space **204** related to the network stack **267**, network traffic or packets received or transmitted by the network stack **267**. In other embodiments, the kernel-level data structure may be used by any of the components or processes **232**, **240**, **234**, **236** and **238** to perform the desired operation of the component or process. In one embodiment, a component **232**, **240**, **234**, **236** and **238** is running in kernel mode **204** when using the kernel-level data structure, while in another embodiment, the component **232**, **240**, **234**, **236** and **238** is running in user mode when using the kernel-level data structure. In some embodiments, the kernel-level data structure may be copied or passed to a second kernel-level data structure, or any desired user-level data structure.

The cache manager **232** may comprise software, hardware or any combination of software and hardware to provide cache access, control and management of any type and form of content, such as objects or dynamically generated objects served by the originating servers **106**. The data, objects or content processed and stored by the cache manager **232** may comprise data in any format, such as a markup language, or communicated via any protocol. In some embodiments, the cache manager **232** duplicates original data stored elsewhere or data previously computed, generated or transmitted, in which the original data may require longer access time to fetch, compute or otherwise obtain relative to reading a cache memory element. Once the data is stored in the cache memory element, future use can be made by accessing the cached copy rather than refetching or recomputing the original data, thereby reducing the access time. In some embodiments, the cache memory element may comprise a data object in memory **264** of device **200**. In other embodiments, the cache memory element may comprise memory having a faster access time than memory **264**. In another embodiment, the cache memory element may comprise any type and form of storage element of the device **200**, such as a portion of a hard disk. In some embodiments, the processing unit **262** may provide cache memory for use by the cache manager **232**. In yet further embodiments, the cache manager **232** may use any portion and combination of memory, storage, or the processing unit for caching data, objects, and other content.

Furthermore, the cache manager **232** includes any logic, functions, rules, or operations to perform any embodiments of the techniques of the appliance **200** described herein. For example, the cache manager **232** includes logic or functionality to invalidate objects based on the expiration of an invalidation time period or upon receipt of an invalidation command from a client **102** or server **106**. In some embodiments,

14

the cache manager **232** may operate as a program, service, process or task executing in the kernel space **204**, and in other embodiments, in the user space **202**. In one embodiment, a first portion of the cache manager **232** executes in the user space **202** while a second portion executes in the kernel space **204**. In some embodiments, the cache manager **232** can comprise any type of general purpose processor (GPP), or any other type of integrated circuit, such as a Field Programmable Gate Array (FPGA), Programmable Logic Device (PLD), or Application Specific Integrated Circuit (ASIC).

The policy engine **236** may include, for example, an intelligent statistical engine or other programmable application(s). In one embodiment, the policy engine **236** provides a configuration mechanism to allow a user to identifying, specify, define or configure a caching policy. Policy engine **236**, in some embodiments, also has access to memory to support data structures such as lookup tables or hash tables to enable user-selected caching policy decisions. In other embodiments, the policy engine **236** may comprise any logic, rules, functions or operations to determine and provide access, control and management of objects, data or content being cached by the appliance **200** in addition to access, control and management of security, network traffic, network access, compression or any other function or operation performed by the appliance **200**. Further examples of specific caching policies are further described herein.

The encryption engine **234** comprises any logic, business rules, functions or operations for handling the processing of any security related protocol, such as SSL or TLS, or any function related thereto. For example, the encryption engine **234** encrypts and decrypts network packets, or any portion thereof, communicated via the appliance **200**. The encryption engine **234** may also setup or establish SSL or TLS connections on behalf of the client **102a-102n**, server **106a-106n**, or appliance **200**. As such, the encryption engine **234** provides offloading and acceleration of SSL processing. In one embodiment, the encryption engine **234** uses a tunneling protocol to provide a virtual private network between a client **102a-102n** and a server **106a-106n**. In some embodiments, the encryption engine **234** is in communication with the Encryption processor **260**. In other embodiments, the encryption engine **234** comprises executable instructions running on the encryption processor **260**.

The multi-protocol compression engine **238** comprises any logic, business rules, function or operations for compressing one or more protocols of a network packet, such as any of the protocols used by the network stack **267** of the device **200**. In one embodiment, multi-protocol compression engine **238** compresses bi-directionally between clients **102a-102n** and servers **106a-106n** any TCP/IP based protocol, including Messaging Application Programming Interface (MAPI) (email), File Transfer Protocol (FTP), HyperText Transfer Protocol (HTTP), Common Internet File System (CIFS) protocol (file transfer), Independent Computing Architecture (ICA) protocol, Remote Desktop Protocol (RDP), Wireless Application Protocol (WAP), Mobile IP protocol, and Voice Over IP (VoIP) protocol. In other embodiments, multi-protocol compression engine **238** provides compression of HyperText Markup Language (HTML) based protocols and in some embodiments, provides compression of any markup languages, such as the Extensible Markup Language (XML). In one embodiment, the multi-protocol compression engine **238** provides compression of any high-performance protocol, such as any protocol designed for appliance **200** to appliance **200** communications. In another embodiment, the multi-protocol compression engine **238** compresses any payload of or any communication using a modified transport control pro-

15

tol, such as Transaction TCP (T/TCP), TCP with selection acknowledgements (TCP-SACK), TCP with large windows (TCP-LW), a congestion prediction protocol such as the TCP-Vegas protocol, and a TCP spoofing protocol.

As such, the multi-protocol compression engine **238** accelerates performance for users accessing applications via desktop clients, e.g., Microsoft Outlook and non-Web thin clients, such as any client launched by popular enterprise applications like Oracle, SAP and Siebel, and even mobile clients, such as the Pocket PC. In some embodiments, the multi-protocol compression engine **238** by executing in the kernel mode **204** and integrating with packet processing engine **240** accessing the network stack **267** is able to compress any of the protocols carried by the TCP/IP protocol, such as any application layer protocol.

High speed layer 2-7 integrated packet engine **240**, also generally referred to as a packet processing engine or packet engine, is responsible for managing the kernel-level processing of packets received and transmitted by appliance **200** via network ports **266**. The high speed layer 2-7 integrated packet engine **240** may comprise a buffer for queuing one or more network packets during processing, such as for receipt of a network packet or transmission of a network packet. Additionally, the high speed layer 2-7 integrated packet engine **240** is in communication with one or more network stacks **267** to send and receive network packets via network ports **266**. The high speed layer 2-7 integrated packet engine **240** works in conjunction with encryption engine **234**, cache manager **232**, policy engine **236** and multi-protocol compression logic **238**. In particular, encryption engine **234** is configured to perform SSL processing of packets, policy engine **236** is configured to perform functions related to traffic management such as request-level content switching and request-level cache redirection, and multi-protocol compression logic **238** is configured to perform functions related to compression and decompression of data.

The high speed layer 2-7 integrated packet engine **240** includes a packet processing timer **242**. In one embodiment, the packet processing timer **242** provides one or more time intervals to trigger the processing of incoming, i.e., received, or outgoing, i.e., transmitted, network packets. In some embodiments, the high speed layer 2-7 integrated packet engine **240** processes network packets responsive to the timer **242**. The packet processing timer **242** provides any type and form of signal to the packet engine **240** to notify, trigger, or communicate a time related event, interval or occurrence. In many embodiments, the packet processing timer **242** operates in the order of milliseconds, such as for example 100 ms, 50 ms or 25 ms. For example, in some embodiments, the packet processing timer **242** provides time intervals or otherwise causes a network packet to be processed by the high speed layer 2-7 integrated packet engine **240** at a 10 ms time interval, while in other embodiments, at a 5 ms time interval, and still yet in further embodiments, as short as a 3, 2, or 1 ms time interval. The high speed layer 2-7 integrated packet engine **240** may be interfaced, integrated or in communication with the encryption engine **234**, cache manager **232**, policy engine **236** and multi-protocol compression engine **238** during operation. As such, any of the logic, functions, or operations of the encryption engine **234**, cache manager **232**, policy engine **236** and multi-protocol compression logic **238** may be performed responsive to the packet processing timer **242** and/or the packet engine **240**. Therefore, any of the logic, functions, or operations of the encryption engine **234**, cache manager **232**, policy engine **236** and multi-protocol compression logic **238** may be performed at the granularity of time intervals provided via the packet processing timer **242**, for

16

example, at a time interval of less than or equal to 10 ms. For example, in one embodiment, the cache manager **232** may perform invalidation of any cached objects responsive to the high speed layer 2-7 integrated packet engine **240** and/or the packet processing timer **242**. In another embodiment, the expiry or invalidation time of a cached object can be set to the same order of granularity as the time interval of the packet processing timer **242**, such as at every 10 ms.

In contrast to kernel space **204**, user space **202** is the memory area or portion of the operating system used by user mode applications or programs otherwise running in user mode. A user mode application may not access kernel space **204** directly and uses service calls in order to access kernel services. As shown in FIG. 2, user space **202** of appliance **200** includes a graphical user interface (GUI) **210**, a command line interface (CLI) **212**, shell services **214**, health monitoring program **216**, and daemon services **218**. GUI **210** and CLI **212** provide a means by which a system administrator or other user can interact with and control the operation of appliance **200**, such as via the operating system of the appliance **200** and either is user space **202** or kernel space **204**. The GUI **210** may be any type and form of graphical user interface and may be presented via text, graphical or otherwise, by any type of program or application, such as a browser. The CLI **212** may be any type and form of command line or text-based interface, such as a command line provided by the operating system. For example, the CLI **212** may comprise a shell, which is a tool to enable users to interact with the operating system. In some embodiments, the CLI **212** may be provided via a bash, csh, tcsh, or ksh type shell. The shell services **214** comprises the programs, services, tasks, processes or executable instructions to support interaction with the appliance **200** or operating system by a user via the GUI **210** and/or CLI **212**.

Health monitoring program **216** is used to monitor, check, report and ensure that network systems are functioning properly and that users are receiving requested content over a network. Health monitoring program **216** comprises one or more programs, services, tasks, processes or executable instructions to provide logic, rules, functions or operations for monitoring any activity of the appliance **200**. In some embodiments, the health monitoring program **216** intercepts and inspects any network traffic passed via the appliance **200**. In other embodiments, the health monitoring program **216** interfaces by any suitable means and/or mechanisms with one or more of the following: the encryption engine **234**, cache manager **232**, policy engine **236**, multi-protocol compression logic **238**, packet engine **240**, daemon services **218**, and shell services **214**. As such, the health monitoring program **216** may call any application programming interface (API) to determine a state, status, or health of any portion of the appliance **200**. For example, the health monitoring program **216** may ping or send a status inquiry on a periodic basis to check if a program, process, service or task is active and currently running. In another example, the health monitoring program **216** may check any status, error or history logs provided by any program, process, service or task to determine any condition, status or error with any portion of the appliance **200**.

Daemon services **218** are programs that run continuously or in the background and handle periodic service requests received by appliance **200**. In some embodiments, a daemon service may forward the requests to other programs or processes, such as another daemon service **218** as appropriate. As known to those skilled in the art, a daemon service **218** may run unattended to perform continuous or periodic system wide functions, such as network control, or to perform any desired task. In some embodiments, one or more daemon

17

services **218** run in the user space **202**, while in other embodiments, one or more daemon services **218** run in the kernel space.

Referring now to FIG. 2B, another embodiment of the appliance **200** is depicted. In brief overview, the appliance **200** provides one or more of the following services, functionality or operations: SSL VPN connectivity **280**, switching/load balancing **284**, Domain Name Service resolution **286**, acceleration **288** and an application firewall **290** for communications between one or more clients **102** and one or more servers **106**. In one embodiment, the appliance **200** comprises any of the network devices manufactured by Citrix Systems, Inc. of Ft. Lauderdale Fla., referred to as Citrix NetScaler devices. Each of the servers **106** may provide one or more network related services **270a-270n** (referred to as services **270**). For example, a server **106** may provide an http service **270**. The appliance **200** comprises one or more virtual servers or virtual internet protocol servers, referred to as a vServer, VIP server, or just VIP **275a-275n** (also referred herein as vServer **275**). The vServer **275** receives, intercepts or otherwise processes communications between a client **102** and a server **106** in accordance with the configuration and operations of the appliance **200**.

The vServer **275** may comprise software, hardware or any combination of software and hardware. The vServer **275** may comprise any type and form of program, service, task, process or executable instructions operating in user mode **202**, kernel mode **204** or any combination thereof in the appliance **200**. The vServer **275** includes any logic, functions, rules, or operations to perform any embodiments of the techniques described herein, such as SSL VPN **280**, switching/load balancing **284**, Domain Name Service resolution **286**, acceleration **288** and an application firewall **290**. In some embodiments, the vServer **275** establishes a connection to a service **270** of a server **106**. The service **275** may comprise any program, application, process, task or set of executable instructions capable of connecting to and communicating to the appliance **200**, client **102** or vServer **275**. For example, the service **275** may comprise a web server, http server, ftp, email or database server. In some embodiments, the service **270** is a daemon process or network driver for listening, receiving and/or sending communications for an application, such as email, database or an enterprise application. In some embodiments, the service **270** may communicate on a specific IP address, or IP address and port.

In some embodiments, the vServer **275** applies one or more policies of the policy engine **236** to network communications between the client **102** and server **106**. In one embodiment, the policies are associated with a VServer **275**. In another embodiment, the policies are based on a user, or a group of users. In yet another embodiment, a policy is global and applies to one or more vServers **275a-275n**, and any user or group of users communicating via the appliance **200**. In some embodiments, the policies of the policy engine have conditions upon which the policy is applied based on any content of the communication, such as internet protocol address, port, protocol type, header or fields in a packet, or the context of the communication, such as user, group of the user, vServer **275**, transport layer connection, and/or identification or attributes of the client **102** or server **106**.

In other embodiments, the appliance **200** communicates or interfaces with the policy engine **236** to determine authentication and/or authorization of a remote user or a remote client **102** to access the computing environment **15**, application, and/or data file from a server **106**. In another embodiment, the appliance **200** communicates or interfaces with the policy engine **236** to determine authentication and/or authorization

18

of a remote user or a remote client **102** to have the application delivery system **190** deliver one or more of the computing environment **15**, application, and/or data file. In yet another embodiment, the appliance **200** establishes a VPN or SSL VPN connection based on the policy engine's **236** authentication and/or authorization of a remote user or a remote client **103**. In one embodiment, the appliance **102** controls the flow of network traffic and communication sessions based on policies of the policy engine **236**. For example, the appliance **200** may control the access to a computing environment **15**, application or data file based on the policy engine **236**.

In some embodiments, the vServer **275** establishes a transport layer connection, such as a TCP or UDP connection with a client **102** via the client agent **120**. In one embodiment, the vServer **275** listens for and receives communications from the client **102**. In other embodiments, the vServer **275** establishes a transport layer connection, such as a TCP or UDP connection with a client server **106**. In one embodiment, the vServer **275** establishes the transport layer connection to an internet protocol address and port of a server **270** running on the server **106**. In another embodiment, the vServer **275** associates a first transport layer connection to a client **102** with a second transport layer connection to the server **106**. In some embodiments, a vServer **275** establishes a pool of transport layer connections to a server **106** and multiplexes client requests via the pooled transport layer connections.

In some embodiments, the appliance **200** provides an SSL VPN connection **280** between a client **102** and a server **106**. For example, a client **102** on a first network **102** requests to establish a connection to a server **106** on a second network **104'**. In some embodiments, the second network **104'** is not routable from the first network **104**. In other embodiments, the client **102** is on a public network **104** and the server **106** is on a private network **104'**, such as a corporate network. In one embodiment, the client agent **120** intercepts communications of the client **102** on the first network **104**, encrypts the communications, and transmits the communications via a first transport layer connection to the appliance **200**. The appliance **200** associates the first transport layer connection on the first network **104** to a second transport layer connection to the server **106** on the second network **104**. The appliance **200** receives the intercepted communication from the client agent **102**, decrypts the communications, and transmits the communication to the server **106** on the second network **104** via the second transport layer connection. The second transport layer connection may be a pooled transport layer connection. As such, the appliance **200** provides an end-to-end secure transport layer connection for the client **102** between the two networks **104**, **104'**.

In one embodiment, the appliance **200** hosts an intranet internet protocol or intranetIP **282** address of the client **102** on the virtual private network **104**. The client **102** has a local network identifier, such as an internet protocol (IP) address and/or host name on the first network **104**. When connected to the second network **104'** via the appliance **200**, the appliance **200** establishes, assigns or otherwise provides an IntranetIP, which is network identifier, such as IP address and/or host name, for the client **102** on the second network **104'**. The appliance **200** listens for and receives on the second or private network **104'** for any communications directed towards the client **102** using the client's established IntranetIP **282**. In one embodiment, the appliance **200** acts as or on behalf of the client **102** on the second private network **104**. For example, in another embodiment, a vServer **275** listens for and responds to communications to the IntranetIP **282** of the client **102**. In some embodiments, if a computing device **100** on the second network **104'** transmits a request, the appliance **200** processes

the request as if it were the client **102**. For example, the appliance **200** may respond to a ping to the client's IntranetIP **282**. In another example, the appliance may establish a connection, such as a TCP or UDP connection, with computing device **100** on the second network **104** requesting a connection with the client's IntranetIP **282**.

In some embodiments, the appliance **200** provides one or more of the following acceleration techniques **288** to communications between the client **102** and server **106**: 1) compression; 2) decompression; 3) Transmission Control Protocol pooling; 4) Transmission Control Protocol multiplexing; 5) Transmission Control Protocol buffering; and 6) caching. In one embodiment, the appliance **200** relieves servers **106** of much of the processing load caused by repeatedly opening and closing transport layer connections to clients **102** by opening one or more transport layer connections with each server **106** and maintaining these connections to allow repeated data accesses by clients via the Internet. This technique is referred to herein as "connection pooling".

In some embodiments, in order to seamlessly splice communications from a client **102** to a server **106** via a pooled transport layer connection, the appliance **200** translates or multiplexes communications by modifying sequence number and acknowledgment numbers at the transport layer protocol level. This is referred to as "connection multiplexing". In some embodiments, no application layer protocol interaction is required. For example, in the case of an in-bound packet (that is, a packet received from a client **102**), the source network address of the packet is changed to that of an output port of appliance **200**, and the destination network address is changed to that of the intended server. In the case of an outbound packet (that is, one received from a server **106**), the source network address is changed from that of the server **106** to that of an output port of appliance **200** and the destination address is changed from that of appliance **200** to that of the requesting client **102**. The sequence numbers and acknowledgment numbers of the packet are also translated to sequence numbers and acknowledgement expected by the client **102** on the appliance's **200** transport layer connection to the client **102**. In some embodiments, the packet checksum of the transport layer protocol is recalculated to account for these translations.

In another embodiment, the appliance **200** provides switching or load-balancing functionality **284** for communications between the client **102** and server **106**. In some embodiments, the appliance **200** distributes traffic and directs client requests to a server **106** based on layer 4 or application-layer request data. In one embodiment, although the network layer or layer 2 of the network packet identifies a destination server **106**, the appliance **200** determines the server **106** to distribute the network packet by application information and data carried as payload of the transport layer packet. In one embodiment, the health monitoring programs **216** of the appliance **200** monitor the health of servers to determine the server **106** for which to distribute a client's request. In some embodiments, if the appliance **200** detects a server **106** is not available or has a load over a predetermined threshold, the appliance **200** can direct or distribute client requests to another server **106**.

In some embodiments, the appliance **200** acts as a Domain Name Service (DNS) resolver or otherwise provides resolution of a DNS request from clients **102**. In some embodiments, the appliance intercepts a DNS request transmitted by the client **102**. In one embodiment, the appliance **200** responds to a client's DNS request with an IP address of or hosted by the appliance **200**. In this embodiment, the client **102** transmits network communication for the domain name

to the appliance **200**. In another embodiment, the appliance **200** responds to a client's DNS request with an IP address of or hosted by a second appliance **200'**. In some embodiments, the appliance **200** responds to a client's DNS request with an IP address of a server **106** determined by the appliance **200**.

In yet another embodiment, the appliance **200** provides application firewall functionality **290** for communications between the client **102** and server **106**. In one embodiment, the policy engine **236** provides rules for detecting and blocking illegitimate requests. In some embodiments, the application firewall **290** protects against denial of service (DoS) attacks. In other embodiments, the appliance inspects the content of intercepted requests to identify and block application-based attacks. In some embodiments, the rules/policy engine **236** comprises one or more application firewall or security control policies for providing protections against various classes and types of web or Internet based vulnerabilities, such as one or more of the following: 1) buffer overflow, 2) CGI-BIN parameter manipulation, 3) form/hidden field manipulation, 4) forceful browsing, 5) cookie or session poisoning, 6) broken access control list (ACLs) or weak passwords, 7) cross-site scripting (XSS), 8) command injection, 9) SQL injection, 10) error triggering sensitive information leak, 11) insecure use of cryptography, 12) server misconfiguration, 13) back doors and debug options, 14) website defacement, 15) platform or operating systems vulnerabilities, and 16) zero-day exploits. In an embodiment, the application firewall **290** provides HTML form field protection in the form of inspecting or analyzing the network communication for one or more of the following: 1) required fields are returned, 2) no added field allowed, 3) read-only and hidden field enforcement, 4) drop-down list and radio button field conformance, and 5) form-field max-length enforcement. In some embodiments, the application firewall **290** ensures cookies are not modified. In other embodiments, the application firewall **290** protects against forceful browsing by enforcing legal URLs.

In still yet other embodiments, the application firewall **290** protects any confidential information contained in the network communication. The application firewall **290** may inspect or analyze any network communication in accordance with the rules or policies of the engine **236** to identify any confidential information in any field of the network packet. In some embodiments, the application firewall **290** identifies in the network communication one or more occurrences of a credit card number, password, social security number, name, patient code, contact information, and age. The encoded portion of the network communication may comprise these occurrences or the confidential information. Based on these occurrences, in one embodiment, the application firewall **290** may take a policy action on the network communication, such as prevent transmission of the network communication. In another embodiment, the application firewall **290** may rewrite, remove or otherwise mask such identified occurrence or confidential information.

Referring now to FIG. 3, an embodiment of the client agent **120** is depicted. The client **102** includes a client agent **120** for establishing and exchanging communications with the appliance **200** and/or server **106** via a network **104**. In brief overview, the client **102** operates on computing device **100** having an operating system with a kernel mode **302** and a user mode **303**, and a network stack **310** with one or more layers **310a-310b**. The client **102** may have installed and/or execute one or more applications. In some embodiments, one or more applications may communicate via the network stack **310** to a network **104**. One of the applications, such as a web browser, may also include a first program **322**. For example, the first

21

program 322 may be used in some embodiments to install and/or execute the client agent 120, or any portion thereof. The client agent 120 includes an interception mechanism, or interceptor 350, for intercepting network communications from the network stack 310 from the one or more applica-

tions. The network stack 310 of the client 102 may comprise any type and form of software, or hardware, or any combinations thereof, for providing connectivity to and communications with a network. In one embodiment, the network stack 310 comprises a software implementation for a network protocol suite. The network stack 310 may comprise one or more network layers, such as any networks layers of the Open Systems Interconnection (OSI) communications model as those skilled in the art recognize and appreciate. As such, the network stack 310 may comprise any type and form of protocols for any of the following layers of the OSI model: 1) physical link layer, 2) data link layer, 3) network layer, 4) transport layer, 5) session layer, 6) presentation layer, and 7) application layer. In one embodiment, the network stack 310 may comprise a transport control protocol (TCP) over the network layer protocol of the internet protocol (IP), generally referred to as TCP/IP. In some embodiments, the TCP/IP protocol may be carried over the Ethernet protocol, which may comprise any of the family of IEEE wide-area-network (WAN) or local-area-network (LAN) protocols, such as those protocols covered by the IEEE 802.3. In some embodiments, the network stack 310 comprises any type and form of a wireless protocol, such as IEEE 802.11 and/or mobile internet protocol.

In view of a TCP/IP based network, any TCP/IP based protocol may be used, including Messaging Application Programming Interface (MAPI) (email), File Transfer Protocol (FTP), HyperText Transfer Protocol (HTTP), Common Internet File System (CIFS) protocol (file transfer), Independent Computing Architecture (ICA) protocol, Remote Desktop Protocol (RDP), Wireless Application Protocol (WAP), Mobile IP protocol, and Voice Over IP (VoIP) protocol. In another embodiment, the network stack 310 comprises any type and form of transport control protocol, such as a modified transport control protocol, for example a Transaction TCP (T/TCP), TCP with selection acknowledgements (TCP-SACK), TCP with large windows (TCP-LW), a congestion prediction protocol such as the TCP-Vegas protocol, and a TCP spoofing protocol. In other embodiments, any type and form of user datagram protocol (UDP), such as UDP over IP, may be used by the network stack 310, such as for voice communications or real-time data communications.

Furthermore, the network stack 310 may include one or more network drivers supporting the one or more layers, such as a TCP driver or a network layer driver. The network drivers may be included as part of the operating system of the computing device 100 or as part of any network interface cards or other network access components of the computing device 100. In some embodiments, any of the network drivers of the network stack 310 may be customized, modified or adapted to provide a custom or modified portion of the network stack 310 in support of any of the techniques described herein. In other embodiments, the acceleration program 120 is designed and constructed to operate with or work in conjunction with the network stack 310 installed or otherwise provided by the operating system of the client 102.

The network stack 310 comprises any type and form of interfaces for receiving, obtaining, providing or otherwise accessing any information and data related to network communications of the client 102. In one embodiment, an interface to the network stack 310 comprises an application pro-

22

gramming interface (API). The interface may also comprise any function call, hooking or filtering mechanism, event or call back mechanism, or any type of interfacing technique. The network stack 310 via the interface may receive or provide any type and form of data structure, such as an object, related to functionality or operation of the network stack 310. For example, the data structure may comprise information and data related to a network packet or one or more network packets. In some embodiments, the data structure comprises a portion of the network packet processed at a protocol layer of the network stack 310, such as a network packet of the transport layer. In some embodiments, the data structure 325 comprises a kernel-level data structure, while in other embodiments, the data structure 325 comprises a user-mode data structure. A kernel-level data structure may comprise a data structure obtained or related to a portion of the network stack 310 operating in kernel-mode 302, or a network driver or other software running in kernel-mode 302, or any data structure obtained or received by a service, process, task, thread or other executable instructions running or operating in kernel-mode of the operating system.

Additionally, some portions of the network stack 310 may execute or operate in kernel-mode 302, for example, the data link or network layer, while other portions execute or operate in user-mode 303, such as an application layer of the network stack 310. For example, a first portion 310a of the network stack may provide user-mode access to the network stack 310 to an application while a second portion 310a of the network stack 310 provides access to a network. In some embodiments, a first portion 310a of the network stack may comprise one or more upper layers of the network stack 310, such as any of layers 5-7. In other embodiments, a second portion 310b of the network stack 310 comprises one or more lower layers, such as any of layers 1-4. Each of the first portion 310a and second portion 310b of the network stack 310 may comprise any portion of the network stack 310, at any one or more network layers, in user-mode 203, kernel-mode, 202, or combinations thereof, or at any portion of a network layer or interface point to a network layer or any portion of or interface point to the user-mode 203 and kernel-mode 203.

The interceptor 350 may comprise software, hardware, or any combination of software and hardware. In one embodiment, the interceptor 350 intercept a network communication at any point in the network stack 310, and redirects or transmits the network communication to a destination desired, managed or controlled by the interceptor 350 or client agent 120. For example, the interceptor 350 may intercept a network communication of a network stack 310 of a first network and transmit the network communication to the appliance 200 for transmission on a second network 104. In some embodiments, the interceptor 350 comprises any type of interceptor 350 comprises a driver, such as a network driver constructed and designed to interface and work with the network stack 310. In some embodiments, the client agent 120 and/or interceptor 350 operates at one or more layers of the network stack 310, such as at the transport layer. In one embodiment, the interceptor 350 comprises a filter driver, hooking mechanism, or any form and type of suitable network driver interface that interfaces to the transport layer of the network stack, such as via the transport driver interface (TDI). In some embodiments, the interceptor 350 interfaces to a first protocol layer, such as the transport layer and another protocol layer, such as any layer above the transport protocol layer, for example, an application protocol layer. In one embodiment, the interceptor 350 may comprise a driver complying with the Network Driver Interface Specification (NDIS), or a NDIS driver. In another embodiment, the interceptor 350 may comprise a

23

min-filter or a mini-port driver. In one embodiment, the interceptor 350, or portion thereof, operates in kernel-mode 202. In another embodiment, the interceptor 350, or portion thereof, operates in user-mode 203. In some embodiments, a portion of the interceptor 350 operates in kernel-mode 202 while another portion of the interceptor 350 operates in user-mode 203. In other embodiments, the client agent 120 operates in user-mode 203 but interfaces via the interceptor 350 to a kernel-mode driver, process, service, task or portion of the operating system, such as to obtain a kernel-level data structure 225. In further embodiments, the interceptor 350 is a user-mode application or program, such as application.

In one embodiment, the interceptor 350 intercepts any transport layer connection requests. In these embodiments, the interceptor 350 execute transport layer application programming interface (API) calls to set the destination information, such as destination IP address and/or port to a desired location for the location. In this manner, the interceptor 350 intercepts and redirects the transport layer connection to a IP address and port controlled or managed by the interceptor 350 or client agent 120. In one embodiment, the interceptor 350 sets the destination information for the connection to a local IP address and port of the client 102 on which the client agent 120 is listening. For example, the client agent 120 may comprise a proxy service listening on a local IP address and port for redirected transport layer communications. In some embodiments, the client agent 120 then communicates the redirected transport layer communication to the appliance 200.

In some embodiments, the interceptor 350 intercepts a Domain Name Service (DNS) request. In one embodiment, the client agent 120 and/or interceptor 350 resolves the DNS request. In another embodiment, the interceptor transmits the intercepted DNS request to the appliance 200 for DNS resolution. In one embodiment, the appliance 200 resolves the DNS request and communicates the DNS response to the client agent 120. In some embodiments, the appliance 200 resolves the DNS request via another appliance 200' or a DNS server 106.

In yet another embodiment, the client agent 120 may comprise two agents 120 and 120'. In one embodiment, a first agent 120 may comprise an interceptor 350 operating at the network layer of the network stack 310. In some embodiments, the first agent 120 intercepts network layer requests such as Internet Control Message Protocol (ICMP) requests (e.g., ping and traceroute). In other embodiments, the second agent 120' may operate at the transport layer and intercept transport layer communications. In some embodiments, the first agent 120 intercepts communications at one layer of the network stack 210 and interfaces with or communicates the intercepted communication to the second agent 120'.

The client agent 120 and/or interceptor 350 may operate at or interface with a protocol layer in a manner transparent to any other protocol layer of the network stack 310. For example, in one embodiment, the interceptor 350 operates or interfaces with the transport layer of the network stack 310 transparently to any protocol layer below the transport layer, such as the network layer, and any protocol layer above the transport layer, such as the session, presentation or application layer protocols. This allows the other protocol layers of the network stack 310 to operate as desired and without modification for using the interceptor 350. As such, the client agent 120 and/or interceptor 350 can interface with the transport layer to secure, optimize, accelerate, route or load-balance any communications provided via any protocol carried by the transport layer, such as any application layer protocol over TCP/IP.

24

Furthermore, the client agent 120 and/or interceptor may operate at or interface with the network stack 310 in a manner transparent to any application, a user of the client 102, and any other computing device, such as a server, in communications with the client 102. The client agent 120 and/or interceptor 350 may be installed and/or executed on the client 102 in a manner without modification of an application. In some embodiments, the user of the client 102 or a computing device in communications with the client 102 are not aware of the existence, execution or operation of the client agent 120 and/or interceptor 350. As such, in some embodiments, the client agent 120 and/or interceptor 350 is installed, executed, and/or operated transparently to an application, user of the client 102, another computing device, such as a server, or any of the protocol layers above and/or below the protocol layer interfaced to by the interceptor 350.

The client agent 120 includes an acceleration program 302, a streaming client 306, and/or a collection agent 304. In one embodiment, the client agent 120 comprises an Independent Computing Architecture (ICA) client, or any portion thereof, developed by Citrix Systems, Inc. of Fort Lauderdale, Fla., and is also referred to as an ICA client. In some embodiments, the client 120 comprises an application streaming client 306 for streaming an application from a server 106 to a client 102. In some embodiments, the client agent 120 comprises an acceleration program 302 for accelerating communications between client 102 and server 106. In another embodiment, the client agent 120 includes a collection agent 304 for performing end-point detection/scanning and collecting end-point information for the appliance 200 and/or server 106.

In some embodiments, the acceleration program 302 comprises a client-side acceleration program for performing one or more acceleration techniques to accelerate, enhance or otherwise improve a client's communications with and/or access to a server 106, such as accessing an application provided by a server 106. The logic, functions, and/or operations of the executable instructions of the acceleration program 302 may perform one or more of the following acceleration techniques: 1) multi-protocol compression, 2) transport control protocol pooling, 3) transport control protocol multiplexing, 4) transport control protocol buffering, and 5) caching via a cache manager. Additionally, the acceleration program 302 may perform encryption and/or decryption of any communications received and/or transmitted by the client 102. In some embodiments, the acceleration program 302 performs one or more of the acceleration techniques in an integrated manner or fashion. Additionally, the acceleration program 302 can perform compression on any of the protocols, or multiple-protocols, carried as payload of network packet of the transport layer protocol. The streaming client 306 comprises an application, program, process, service, task or executable instructions for receiving and executing a streamed application from a server 106. A server 106 may stream one or more application data files to the streaming client 306 for playing, executing or otherwise causing to be executed the application on the client 102. In some embodiments, the server 106 transmits a set of compressed or packaged application data files to the streaming client 306. In some embodiments, the plurality of application files are compressed and stored on a file server within an archive file such as a CAB, ZIP, SIT, TAR, JAR or other archive. In one embodiment, the server 106 decompresses, unpackages or unarchives the application files and transmits the files to the client 102. In another embodiment, the client 102 decompresses, unpackages or unarchives the application files. The streaming client 306 dynamically installs the application, or portion thereof, and executes the application. In one embodiment, the streaming client 306

25

may be an executable program. In some embodiments, the streaming client **306** may be able to launch another executable program.

The collection agent **304** comprises an application, program, process, service, task or executable instructions for identifying, obtaining and/or collecting information about the client **102**. In some embodiments, the appliance **200** transmits the collection agent **304** to the client **102** or client agent **120**. The collection agent **304** may be configured according to one or more policies of the policy engine **236** of the appliance. In other embodiments, the collection agent **304** transmits collected information on the client **102** to the appliance **200**. In one embodiment, the policy engine **236** of the appliance **200** uses the collected information to determine and provide access, authentication and authorization control of the client's connection to a network **104**.

In one embodiment, the collection agent **304** comprises an end-point detection and scanning mechanism, which identifies and determines one or more attributes or characteristics of the client. For example, the collection agent **304** may identify and determine any one or more of the following client-side attributes: 1) the operating system and/or a version of an operating system, 2) a service pack of the operating system, 3) a running service, 4) a running process, and 5) a file. The collection agent **304** may also identify and determine the presence or versions of any one or more of the following on the client: 1) antivirus software, 2) personal firewall software, 3) anti-spam software, and 4) internet security software. The policy engine **236** may have one or more policies based on any one or more of the attributes or characteristics of the client or client-side attributes.

In some embodiments and still referring to FIG. 3, a first program **322** may be used to install and/or execute the client agent **120**, or portion thereof, such as the interceptor **350**, automatically, silently, transparently, or otherwise. In one embodiment, the first program **322** comprises a plugin component, such as an ActiveX control or Java control or script that is loaded into and executed by an application. For example, the first program comprises an ActiveX control loaded and run by a web browser application, such as in the memory space or context of the application. In another embodiment, the first program **322** comprises a set of executable instructions loaded into and run by the application, such as a browser. In one embodiment, the first program **322** comprises a designed and constructed program to install the client agent **120**. In some embodiments, the first program **322** obtains, downloads, or receives the client agent **120** via the network from another computing device. In another embodiment, the first program **322** is an installer program or a plug and play manager for installing programs, such as network drivers, on the operating system of the client **102**.

Referring now to FIG. 4, a block diagram depicts one embodiment of a system for routing packets between clients, the system comprising a secure sockets layer virtual private network (an SSL/VPN) network appliance. The SSL/VPN network appliance **200** comprises a receiver **420**, a packet inspector **430** and a transceiver **440**. The receiver **420** receives a packet from a first client **102a** over a connection. The packet inspector **430** identifies, responsive to an inspection of the received packet, 1) a type of connection required for transmission of the received packet to a destination address identified by the received packet and 2) a second client **102b** connected via an SSL/VPN connection to the SSL/VPN network appliance **200** and associated with the identified destination address. The transceiver **440** requests the establish-

26

ment of a connection of the identified type and transmits the received packet to the second client **102b** over the requested connection.

The first client **102a** transmits a packet to the SSL/VPN network appliance **200**. In one embodiment, the first client **102a** establishes a secure sockets layer virtual private network (an SSL/VPN) connection to the SSL/VPN network appliance **200**. In another embodiment, the first client **102a** transmits the packet to the SSL/VPN network appliance **200** over the SSL/VPN connection. In still another embodiment, a virtual server **275** on the SSL/VPN network appliance **200** establishes a transport layer connection, such as a TCP or UDP connection, with a client **102a** and listens for and receives communications from the client **102a**, as described above in connection with FIG. 2B.

In one embodiment, the second client **102b** establishes an initial SSL/VPN connection to the SSL/VPN network appliance **200**. In another embodiment, the second client **102b** receives a request to establish the connection of the identified type from the SSL/VPN network appliance **200** over the initial SSL/VPN connection. In still another embodiment, the second client **102b** establishes a connection of the identified type between the second client **102b** and the SSL/VPN network appliance **200**. In yet another embodiment, the second client **102b** establishes the connection of the identified type within the initial SSL/VPN connection to the SSL/VPN network appliance **200**.

In some embodiments, the SSL/VPN network appliance **200** provides an SSL/VPN connection between the client **102a** and the client **102b**. In one of these embodiments, the client **102a** resides on a different network **104** than the client **102b**, which may reside, for example on network **104'** or a third network **104''**. In another of these embodiments, the SSL/VPN network appliance **200** provides the connection between the client **102a** and the client **102b** using functionality similar to that described in connection with FIG. 2B and the connection provided between a client **102** and a server **106**. In still another of these embodiments, where the first client **102a** and the second client **102b** are connected to the SSL/VPN network appliance **200** via a first and a second transport layer connection, respectively, the SSL/VPN network appliance **200** requests the establishment, by the second client **102b**, of a third transport layer connection. In one of these embodiments, the third transport layer connection is between the SSL/VPN network appliance **200** and the client **102b**. In still another of these embodiments, the appliance **200** provides an end-to-end secure transport layer connection for the client **102a** and the client **102b** between the networks **104** and **104'** or **104''**.

The receiver **420** on the SSL/VPN network appliance **200** receives a packet from a first client **102a** over a connection. In one embodiment, the connection is an SSL/VPN connection from the first client **102a** to the SSL/VPN network appliance **200**. In another embodiment, the connection is the first transport layer connection established between the client **102a** and the SSL/VPN network appliance **200**. In some embodiments, the client agent **120** intercepts communications of the client **102a** on the first network **104**, encrypts the communications, and transmits the communications via a first transport layer connection to the SSL/VPN network appliance **200**.

The packet inspector **430** identifies, responsive to an inspection of the received packet, 1) a type of connection required for transmission of the received packet to a destination address identified by the received packet and 2) a second client **102b** connected via an SSL/VPN connection to the SSL/VPN network appliance **200** and associated with the identified destination address. In one embodiment, the packet

27

is addressed to the second client **102b**. In another embodiment, the destination address is a network address of the second client **102b** on the virtual private network.

In one embodiment, the SSL/VPN network appliance **200** inspects an intercepted packet and identifies the destination address as an address associated with a second client **102b** connected to the SSL/VPN network appliance **200** via a second transport layer connection. In another embodiment, the SSL/VPN network appliance **200** identifies a type of connection required for transmission of the intercepted packet to the second client **102b**. In still another embodiment, the SSL/VPN network appliance **200** determines that the type of connection requires the establishment of a third transport layer connection, for example, because transmission of the packet requires a dedicated connection. The transceiver **440** requests the establishment of a connection of the identified type and transmits the received packet to the second client **102b** over the requested connection.

Referring now to FIG. 5, a flow diagram depicts one embodiment of the steps taken in a method for routing packets between clients. A packet is received from a first client **102a** connected to a secure sockets layer virtual private network (an SSL/VPN) network appliance **200** (step **502**). An identification is made, responsive to an inspection of the received packet, of i) a type of connection required for transmission of the received packet to a destination address identified by the received packet and ii) a second client **102b** connected via an SSL/VPN connection to the SSL/VPN network appliance **200** and associated with the identified destination address (step **504**). A request is made for establishment, by the second client **102b**, of a connection of the identified type within the SSL/VPN connection (step **506**). The received packet is transmitted to the second client **102b** via the established connection of the identified type (step **508**).

A packet is received from a first client **102a** connected to an SSL/VPN network appliance **200** (step **502**). In one embodiment, the SSL/VPN network appliance **200** receives the packet via an SSL/VPN connection. In another embodiment, the client **102a** transmits the packet after establishing an SSL/VPN connection with the SSL/VPN network appliance **200**. In some embodiments, a client **102a** transmits a packet requesting initiation of a TCP connection to client **102b**. In one of these embodiments, the packet is intercepted and transmitted to the client application **120**. In another of these embodiments, a transport data interface (TDI) layer filter intercepts the packet. In still another of these embodiments, the client application **120** initiates a new connection with the SSL/VPN network appliance **200**. In yet another of these embodiments, the packet received from the first client **102a** is a request packet sent by the client application **120** to the SSL/VPN network appliance **200** over the new connection.

An identification is made, responsive to an inspection of the received packet, of i) a type of connection required for transmission of the received packet to a destination address identified by the received packet and ii) a second client **102b** connected via an SSL/VPN connection to the SSL/VPN network appliance **200** and associated with the identified destination address (step **504**). In one embodiment, the SSL/VPN network appliance **200** inspects the received packet and makes the identifications. In another embodiment, the SSL/VPN network appliance **200** identifies the type of connection required as a TCP connection. In still another embodiment, the SSL/VPN network appliance identifies the type of connection required as a UDP connection. In yet another embodiment, the SSL/VPN network appliance **200** determines that the type of connection requires the establishment of a third

28

transport layer connection, for example, because transmission of the packet requires a dedicated connection.

In one embodiment, the SSL/VPN network appliance **200** identifies the destination address as an address assigned to a client **102b** connected to the SSL/VPN network appliance **200**. In another embodiment, the SSL/VPN network appliance **200** identifies the destination address as a hosted intranetIP **282** address of the client **102b** on the virtual private network **104'**. In some embodiments, the SSL/VPN network appliance **200** retrieves an identifier for the client **102b** responsive to a destination intranetIP address **282** contained in the received packet.

A request is made for establishment, by the second client **102b**, of a connection of the identified type within the SSL/VPN connection (step **506**). In one embodiment, the SSL/VPN network appliance **200** requests the establishment, by the second client **102b**, of the connection of the identified type. In another embodiment, the SSL/VPN network appliance **200** requests that the second client **102b** establish a connection of the identified type between the second client **102b** and the SSL/VPN network appliance **200**. In still another embodiment, the second client **102b** establishes a connection of the identified type. In some embodiments, the SSL/VPN network appliance **200** stores the packet addressed to the client **102b** in a queue until the requested connection is established.

In some embodiments, such as those described above in connection with FIG. 2B, the SSL/VPN network appliance **200** responds to communication from a computing device **100** on behalf of a client device **102**. In other embodiments, the SSL/VPN network appliance **200** implements techniques for establishing a server-initiated connection to the client **102b**.

In still other embodiments, such as those depicted by FIG. 5, the SSL/VPN network appliance **200** requests establishment by the client **102b** of a connection between the SSL/VPN network appliance **200** and the client **102b** and transmits over the established connection the packets received from the client **102a**. In one of these embodiments, the SSL/VPN network appliance **200** sends a control packet to the client **102b** over an established control channel. In another of these embodiments, the SSL/VPN network appliance **200** transmits information to the client **102b** including an IP address of the server, an identification of a port on the SSL/VPN network appliance **200**, a client IP address, and a client port identifier. In still another of these embodiments, the information sent to the client is intercepted by a kernel-mode filter, such as a transport data interface (TDI) filter. In yet another of these embodiments, the client agent **120** on the client **102b** establishes the requested connection with a virtual server **275** on the SSL/VPN network appliance **200**. In a further one of these embodiments, the client **102b** transmits a control packet to the virtual server **275** on the SSL/VPN network appliance **200** over the newly-established connection. The control packet may include information received from the SSL/VPN network appliance **200** in the establishment request, such as the IP address of the server, an identification of a port on the SSL/VPN network appliance **200**, a client IP address, and a client port identifier.

The received packet is transmitted to the second client via the established connection of the identified type (step **508**). As described above in connection with FIG. 2B, the SSL/VPN network appliance **200** may listen for and receive on the second or private network **104'** any communications directed towards the client **102b** using the established IntranetIP **282**, and may, in some embodiments, establish a connection, such as a TCP or UDP connection, with computing device **100** on

the second network **104** requesting a connection with the client's IntranetIP **282**. In some embodiments, instead of responding to the computing device **100** on behalf of the client device **102b**, the SSL/VPN network appliance **200** forwards the communication to the client **102b** over a new connection. In one of these embodiments, the SSL/VPN network appliance **200** receives packets from the client **102b** responding to the packets received from the client **102a** via the SSL/VPN network appliance **200**. In another of these embodiments, the SSL/VPN network appliance **200** forwards a response packet to the client **102a**. In still another of these embodiments, the SSL/VPN network appliance **200** uses the connections between itself and the client **102a** and the connection between itself and the client **102b** to provide a peer-to-peer communications channel between the two clients **102a, b** over the virtual private network.

Referring now to FIG. 6, a flow diagram depicts another embodiment of the steps taken in a method for routing packets between clients. A client **102a** initiates a connection to another client **102b** (step **601**). A client agent **120** initiates a new connection to an SSL/VPN network appliance **200** (step **602**). The SSL/VPN network appliance **200** identifies the connection as a connection associated with a first packet control block (step **603**). The client agent **120** transmits a request packet to the SSL/VPN network appliance **200**, requesting a connection to the client **102b** (step **604**). The SSL/VPN network appliance **200** retrieves a session identifier associated with a client agent **120'** of the client **102b** (step **605**). The SSL/VPN network appliance **200** transmits a control packet to the client agent **120'** via an existing connection (step **606**). The client agent **120'** initiates a connection to the client **102b** responsive to information in the control packet (step **607**). A transport data interface (TDI) layer filter intercepts packets sent from the client **102b** and addressed to the server (step **608**). The client agent **120'** establishes a second connection to the SSL/VPN network appliance **200** (step **609**). The client agent **120'** transmits a control packet to the SSL/VPN network appliance **200** via the second connection (step **610**). The SSL/VPN network appliance **200** associates the first packet control block with a second patent control block, the second patent control block identified responsive to information in the received control packet (step **611**).

A client **102a** initiates a connection to a client **102b** (step **601**). In one embodiment an application on the client **102a** initiates the connection. In another embodiment, the client **102a** initiates a TCP connection to the client **102b**. In still another embodiment, the client **102a** addresses a packet containing a request for initiation of the connection to the intranetIP address of the client **102b**. In yet another embodiment, a packet sent from the client **102a** and addressed to the intranetIP address of the client **102b** is intercepted. In some embodiments, a TDI layer filter intercepts the packet.

A client agent **120** initiates a new connection to an SSL/VPN network appliance **200** (step **602**). In one embodiment, the client agent **120** transmits a packet to the SSL/VPN network appliance **200** requesting initiation of the new connection. In another embodiment, the client agent **120** transmits a packet initiating the new connection across an existing UDP connection. In still another embodiment, the client agent **120** transmits information including, but not limited to, an identification a type of connection requested (such as TCP or UDP), an intranetIP address of the client **102b**, a port, an identification of a virtual server on the SSL/VPN network appliance **200**, and a port on the SSL/VPN network appliance **200**.

The SSL/VPN network appliance **200** identifies the connection as a connection associated with a first packet control

block (step **603**). In some embodiments, control packets received from the client agent **120** are addressed to or associated with the first packet control block.

The client agent **120** transmits a request packet to the SSL/VPN network appliance **200**, requesting a connection to the client **102b** (step **604**). In one embodiment, the client agent **120** transmits the request over a UDP control channel. In another embodiment, the request packet includes the intranetIP address of the client **102b**.

The SSL/VPN network appliance **200** retrieves a session identifier associated with a client agent **120'** of the client **102b** (step **605**). In one embodiment, the SSL/VPN network appliance **200** retrieves the session identifier associated with the client agent **120'** responsive to a destination intranetIP address **282** contained in the received packet. In another embodiment, the SSL/VPN network appliance **200** associates the session identifier with a queue of packets to be transmitted to the client **102b**. In still another embodiment, the SSL/VPN network appliance **200** stores a received packet addressed to the intranetIP address of the client **102b** in the queue associated with the session identifier. In other embodiments, the first patent control block is associated with the session identifier.

The SSL/VPN network appliance **200** transmits a control packet to the client agent **120'** via an existing connection (step **606**). In some embodiments, the control packet includes information such as, but not limited to, an IP address of the server, an identification of a port on the SSL/VPN network appliance **200**, a client IP address, and a client port identifier. In other embodiments, the control packet is addressed to the client **102b**. In one of these embodiments, a TDI layer filter intercepts the packet and forwards it to the client agent **120'**.

The client agent **120'** initiates a connection to the client **102b** responsive to information in the control packet (step **607**). In one embodiment, the client agent **120'** forwards packets to the client **102b** received from the client **102a** via the SSL/VPN network appliance **200**.

A transport data interface (TDI) layer filter intercepts packets sent from the client **102b** and addressed to the server (step **608**). In one embodiment, the TDI layer filter forwards the intercepted packets to the client agent **120'**.

The client agent **120'** establishes a second connection to the SSL/VPN network appliance **200** (step **609**). In one embodiment, the client agent **120'** establishes the second connection to a virtual server **275** on the SSL/VPN network appliance **200**. In another embodiment, the SSL/VPN network appliance **200** associates a second packet control block with the established connection.

The client agent **120'** transmits a control packet to the SSL/VPN network appliance **200** via the second connection (step **610**). The control packet may include information received from the SSL/VPN network appliance **200** in the establishment request, such as the IP address of the server, an identification of a port on the SSL/VPN network appliance **200**, a client IP address, and a client port identifier.

The SSL/VPN network appliance **200** associates the first packet control block with a second patent control block, the second packet control block identified responsive to information in the received control packet (step **611**). In one embodiment, the second packet control block is associated with the connection established by the client agent **120'** to the SSL/VPN network appliance **200**. In some embodiments, the SSL/VPN network appliance **200** receives packets from the client **102a** over the connection associated with the first control packet block and forwards them to the client **102b** using the established connection associated with the second packet control block. In other embodiments, the SSL/VPN network

31

appliance **200** receives packets from the client **102b** addressed to an intranetIP address of the client **102a** over the established connection and forwards the packets to the client **102** over the connection associated with the first control packet block. In one embodiment, the SSL/VPN network appliance **200** uses the established connections and associated control packet blocks to forward packets between the client **102a** and the client **102b**, providing a peer-to-peer communications channel between the two clients **102a**, **b** over the virtual private network.

The present invention may be provided as one or more computer-readable programs embodied on or in one or more articles of manufacture. The article of manufacture may be a floppy disk, a hard disk, a CD-ROM, a flash memory card, a PROM, a RAM, a ROM, or a magnetic tape. In general, the computer-readable programs may be implemented in any programming language, LISP, PERL, C, C++, PROLOG, or any byte code language such as JAVA. The software programs may be stored on or in one or more articles of manufacture as object code.

Having described certain embodiments of the invention, it will now become apparent to one of skill in the art that other embodiments incorporating the concepts of the invention may be used. Therefore, the invention should not be limited to certain embodiments, but rather should be limited only by the spirit and scope of the following claims.

What is claimed is:

1. A method for routing packets between clients on a virtual private network, the method comprising:

- (a) identifying, by a device, from a packet of a first application of a first client of a plurality of clients received via a first virtual private network connection between the first client and the device, a second client corresponding to a destination of the packet and a type of transport layer connection to communicate with the second client,
- (b) establishing, by the device, via a second virtual private network connection determined by the device to be already established between the device and the second client corresponding to the destination, a transport layer connection with an agent on the second client that corresponds to the type of transport layer connection, to provide a peer-to-peer channel between the first client and the second client using the established transport layer connection; and
- (c) transmitting, by the device, the packet via the transport layer connection within the second virtual private network connection to the agent on the second client for communicating the packet to a second application on the second client.

2. The method of claim 1, further comprising communicating, by the agent, the packet to the second application via a second transport layer connection established between the agent and the second application.

3. The method of claim 1, wherein step (a) further comprises inspecting, by the device, the packet to determine the type of connection and the destination.

4. The method of claim 1, wherein step (a) further comprises determining, by the device, that the second client corresponding to the destination has the second virtual private network connection with the device.

5. The method of claim 1, wherein step (a) further comprises receiving, by the device, the packet from a second agent of the first client, the second agent intercepting the packet transmitted by the first application to the second application.

6. The method of claim 1, wherein the first virtual private network connection and the second virtual private network connection comprises a secure socket layer.

32

7. The method of claim 1, wherein the type of transport layer connection comprises a transport control protocol (TCP).

8. The method of claim 1, wherein the type of transport layer connection comprises a user datagram protocol (UDP).

9. The method of claim 1, wherein step (b) further comprises communicating, by the device, via the second virtual private network connection, a request to the agent to establish the transport layer connection of the type of transport layer connection with the device.

10. The method of claim 1, further comprising the first application of the first client and the second application of the second client communicating over the secure virtual private network connection and the second virtual private network connection using the transport layer connection of the identified type.

11. A system for routing packets between clients on a virtual private network, the system comprising:

- a circuitry configured to identify from a packet of a first application of a first client of a plurality of clients received via a first virtual private network connection between the first client and the circuitry, a second client corresponding to a destination of the packet and a type of transport layer connection to communicate with the second client,
- wherein the circuitry is configured to establish via a second virtual private network connection determined by the circuitry to be already established between the circuitry and the second client corresponding to the destination, a transport layer connection with an agent on the second client that corresponds to the type of transport layer connection, to provide a peer-to-peer channel between the first client and the second client using the established transport layer connection; and
- wherein the circuitry is configured to transmit the packet via the transport layer connection within the second virtual private network connection to the agent on the second client for communicating the packet to a second application on the second client.

12. The system of claim 11, wherein the agent is further configured to communicate the packet to the second application via a second transport layer connection established between the agent and the second application.

13. The system of claim 11, wherein the circuitry is further configured to inspect the packet to determine the type of connection and the destination.

14. The system of claim 11, wherein the agent is further configured to determine that the second client corresponding to the destination has the second virtual private network connection with the circuitry.

15. The system of claim 11, wherein the agent is further configured to receive the packet from a second agent of the first client, the second agent intercepting the packet transmitted by the first application to the second application.

16. The system of claim 11, wherein the first virtual private network connection and the second virtual private network connection comprises a secure socket layer.

17. The system of claim 11, wherein the type of transport layer connection comprises a transport control protocol (TCP).

18. The system of claim 11, wherein the type of transport layer connection comprises a user datagram protocol (UDP).

19. The system of claim 11, wherein the agent is further configured to communicate via the second virtual private network connection, a request to the agent to establish the transport layer connection of the type of transport layer connection with the circuitry.

20. The system of claim 11, wherein the circuitry is configured to have the first application of the first client and the second application of the second client communicate over the secure virtual private network connection and the second virtual private network connection using the transport layer connection of the identified type.

* * * * *